Bereskin & Parr
INTELLECTUAL PROPERTY LAW

November 9, 2004

Kendrick Lo B.A.Sc. (Eng. Sci.), MBA, LL.B.
416 957 1685 klo@bereskinparr.com

Your Reference: 09/718,500
Our Reference: 11483-80

CERTIFIED COPY

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

Re: United States Patent Application Serial No.: 09/718,500
Applicant: DEMBO, Ron et al.
Title: SYSTEM AND METHOD FOR TRADING OFF PUT AND CALL VALUES OF A PORTFOLIO (was "Method of Portfolio Valuation")
Filed: November 24, 2000
Group: 3625

On the filing of this application, the applicant claimed convention priority from Canadian Patent Application No. 2,290,888 filed on November 26, 1999. At this time, the applicant encloses a certified copy of the said Canadian patent application. Acknowledgement of the completion of the claim for convention priority by the filing of the certified copy of the Canadian application is respectfully requested.

Respectfully submitted,

Kendrick Lo
Registration No. 54,948
/sc
Encl.

BEST AVAILABLE COPY

Office de la propriété intellectuelle du Canada

Un organisme d'Industrie Canada

Canadian Intellectual Property Office

An Agency of Industry Canada

Bureau canadien des brevets

Certification

Canadian Patent Office

Certification

La présente atteste que les documents ci-joints, dont la liste figure ci-dessous, sont des copies authentiques des documents déposés au Bureau des brevets.

This is to certify that the documents attached hereto and identified below are true copies of the documents on file in the Patent Office.

Specification and Drawings, as originally filed, with Application for Patent Serial No: **2,290,888,** on November 26, 1999, by **ALGORITHMICS INTERNATIONAL CORP.,** assignee of Andre Aziz, Ron Dembo, Ben Deprisco and Helmut Mausser, for "Risk Management, Pricing and Portfolio Makeup System and Method".

CERTIFIED COPY OF PRIORITY DOCUMENT

Agent certificateur/Certifying Officer

October 7, 2004

Date

Canada

(CIPO 68) 04-09-02

OPIC CIPO

Risk Management, Pricing and Portfolio Makeup System and Method

FIELD OF THE INVENTION

The present invention relates to a method of managing the risk associated with a set of investments and/or financial instruments. More specifically, the present invention relates to a system and method of determining, quantifying and/or managing risk versus return in an investment portfolio and/or determining pricing for investments and/or instruments.

BACKGROUND OF THE INVENTION

Risk Management systems are known and are commonly employed by financial institutions, resource-based corporations, trading organizations, governments or other users to make informed decisions to assess and/or manage the risk associated with the operations of the user.

One popular example of a known risk management system is the RiskWatch V3.1.2 system, sold by the assignee of the present invention. This system is very flexible and allows users to employ models of the instruments in the user's portfolio, which models are evaluated at appropriate time intervals, in view of a range of possible scenarios. Each scenario comprises a set of values for the risk factors employed in the models, at each time interval, and each scenario has an assigned probability. The resulting risk values of the instruments when evaluated under each scenario at each time interval of interest are then used to produce one or more risk metrics which are examined to assess the risk to the user of holding the portfolio of instruments under the evaluated scenarios. Perhaps the most common risk value is the monetary value of the instrument or instruments under consideration, although other risk values including deltas, gammas and other computed values can also be employed. By combining these risk values appropriately, desired risk metrics can be obtained so that the user can identify opportunities for changing the composition of the portfolio to reduce the overall risk or to achieve an acceptable level of risk.

However, problems exist with known risk management systems and methods. For example, while the existence of a general trade off between risk and return is intuitive, in complex trading environments and/or large portfolios it can be difficult to assess and/or to exploit this trade off and yet it is always desired to maximize the return of an investment for selected level of risk. Previous attempts to maximize return for a given level of risk have included using a mean-variance framework or using the Markowitz efficient frontier, which is produced by minimizing portfolio variance subject to specified levels of portfolio return. With these measures, the make up of the

portfolio can be altered to bring the risk/reward measure into line with the desired results.

While such previous attempts have provided a useful assessment of the risk reward trade off, their accuracy and/or merit has been the subject of some debate as these assessments are based on historical information respecting the portfolio. In other words, with existing assessment systems, the risks and rewards are assessed in view of the past performance of the instruments in the portfolio, and thus assumes, explicitly or implicitly, similar results in the future, and any error in this assumption affects the accuracy of the results. Further, such prior attempts have ignored issues such as the aging of investments (i.e. bond coupons maturing into cash, instruments themselves maturing, etc.) and settlement/liquidity of instruments.

Another problem with the prior art risk management systems and methods is that broad simplifying assumptions must be made for instruments or exposures which do not exist (will not be created until some point in the future, eg. - 90 day T-Bill whose start date is two years away) or for which appropriate pricing information is not available. To date, risk management systems have required simplifying assumptions to be made for such circumstances, even when such assumptions may be in conflict with conditions under one or more scenarios.

It is therefore desired to have a method and system of assessing and/or exploiting the trade off between risk and reward to select the make up of a trading portfolio and/or to provide pricing information.

SUMMARY OF THE INVENTION

It is an object of the present invention to provide a novel method and system for risk management, pricing and portfolio makeup which obviates or mitigates at least some of the above-identified disadvantages of the prior art.

DETAILED DESCRIPTION OF THE INVENTION

The assignee of the present invention has developed a new system, framework and methodology for determining and analyzing risk. This new development is referred to herein as a Mark to Future framework and is described in co-pending U.S. patent application 09/323,680, the contents of which are incorporated herein by reference.

The assignee of the present invention has also developed a system and method for use with the Mark to Future framework which allows for the evolution, or aging, of portfolios according to rule-based strategies. This system and method is described in co-pending U.S. patent

application 09/324,920, the contents of which are incorporated herein by reference.

mark to future

mark to future
from Algorithmics

Ron S. Dembo

Imagine an archer capable of hitting a 'mark' as far off as 800 yards. That 'moment of impact' will actually happen in the future. The target is selected, allowances are made for field conditions, target movement, wind velocity and direction, angle of flight, the degree of draw, timing and breathing—and then, *thhhhp!*

For 20,000 years, on every inhabited continent except Australia, the longbow was an extraordinarily effective long distance killer—*a tool to mark the future.* English archers, often outnumbered by as much as 8 to 1, were capable of striking targets at a ratio of 1,000 to 1! Prince Napoleon put it this way, "A first-rate English archer, who in a single minute was unable to draw and discharge his bow 12 times with a range of 240 yards, and in these 12 shots once missed his man, was very lightly esteemed."

mark

This comprehensive guide is the first in a series of documents that will describe the Mark-to-Future methodology in detail.

In Part 1: Overview, we provide a description of the framework advocated by U.S. regulators. This, in conjunction with a brief history of risk management systems and regulatory and supervisory requirements, serves as an introduction for the Mark-to-Future framework. A brief overview of the methodology is provided and concludes with a detailed presentation and discussion of benefits of the methodology.

In Part 2: The Mark-to-Future Methodology, we present a detailed, step-by-step discussion of the methodology. Those interested in understanding the particulars of instrument, product and portfolio mapping will find the pertinent details here. Finally, a new model for trading-off risk and reward, the Put/Call efficient frontier, is described. Also presented is a discussion of the richness of scenario modelling and an introduction to innovative methods of generating forward-looking scenarios. It is our hope that this will stimulate interest in this critical area of risk management research. Anyone concerned about the ability of the Mark-to-Future framework to provide consistent and comprehensive measures of risk and reward and to accommodate changes to practise and supervision should review the selected, not exhaustive, list of measures that can be produced based on the Mark-to-Future Cube.

In Part 3: Technology Considerations, we use a case study based on a hypothetical global bank to illustrate the need for a framework such as Mark-to-Future. The many benefits of the distributed Mark-to-Future process are compared to the limitations of traditional, centralized risk management systems. A case study based on an actual implementation of Mark-to-Future at HypoVereinsbank is used to discuss practical considerations.

In Part 4: Advanced Applications, we present an interesting selection of advanced applications that address important problems in risk management. Topics include: the accurate measurement of counterparty risk, the integration of market and credit risk, the measurement of risk over long horizons using dynamic portfolios and the development of tools to optimally restructure portfolios. Each of these applications is built on a Mark-to-Future Cube.

to future™ contents

CA 02290888 1999-11-26

mark

part 2
The Mark-to-Future Methodology

to future™ contents

mark

part 4
Advanced Mark-to-Future Applications

to future™

www.mark-to-future.com

editors
Judith M. Farvolden, Marianne C. Kupina

authors
Ron S. Dembo, Andrew Aziz, Bob Boettcher,
Judith M. Farvolden, Alexandra Shaw, Michael Zerbs

contributors
Liane Beam, Stephanie Collins, Helmut Mausser,
Tammy McCausland, Ben de Prisco, Dan Rosen

creative & art direction
Scott Thornley + Company Inc.

web site creative
STC WebWorks

web site development
STC WebWorks, Dale Fallon, Marianne C. Kupina

Algorithmics Incorporated

Algorithmics Incorporated was founded in 1989 as a response to the complex issues surrounding financial risk management for the enterprise. Today, as the leading software provider with the most experienced team in the industry, Algorithmics continues to focus its efforts on creating and implementing enterprise risk management software that meets the evolving needs of its customers. Headquartered in Toronto, with 14 offices around the world, Algorithmics serves more than 100 global financial institutions and corporations.

Algorithmics Incorporated
Corporate Headquarters
185 Spadina Avenue,
Toronto, ON Canada M5T 2C6
phone: 416-217-1500
fax: 416-971-6100

info@mark-to-future.com
www.algorithmics.com

What is mark to future?

An Executive Summary by Ron Dembo

Ten years ago, Algorithmics was formed in response to the complex issues surrounding risk management in banking. At that time, even the most advanced financial institutions associated risk management with delta-hedging an options book. We recognized then that to proactively measure and manage risk and reward, one needed to know the total exposure of the institution across all of its global activities. This gave birth to what is today referred to as enterprise risk management or ERM. Since then, our software has helped to transform the way in which approximately 100 banks, asset managers, insurance companies and corporations, in 18 countries, are able to measure their risk and manage their capital. At the heart of our software solution is the Mark-to-Future methodology.

Today, we are making Mark-to-Future generally available as a proposed new standard for risk/reward management. This radically new approach to risk management significantly extends current methodologies. It conforms not only with existing regulatory requirements, but also to proposed regulatory changes. It responds perfectly to the calls by regulators for a comprehensive framework and not 'formulaic approaches to risk' (Greenspan 1999). As the first truly forward-looking risk/reward framework, we believe Mark-to-Future will revolutionize the industry by profoundly affecting the way risk management and capital allocation is conducted in financial institutions. Most importantly, since Mark-to-Future is a generic framework and not a measure, we believe it will be able to evolve with risk management practice.

Mark-to-Future allows for portfolios that may change over time and under differing scenarios. By taking into account the effects of changing portfolios at future points in time, a more realistic assessment of risk is possible. Mark-to-Future also provides a natural basis for linking market, credit, liquidity, legal and other sources of risk. It also provides a unified framework for calculating the risk/reward trade-off of any combinations of these risks.

In a Mark-to-Future world, the focus is on scenarios and instruments, not portfolios. Since the framework is additive, the Mark-to-Future of a portfolio is simply the same combination of each instrument's Mark-to-Future. Accordingly, once a decision has been made with respect to the future scenarios and the instruments one is likely to trade, the Mark-to-Future values of these instruments may be computed well before the actual portfolio

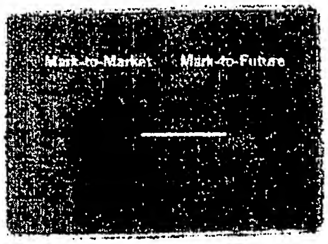

Determine the value of your position today and then choose the horizon over which risk is to be measured.

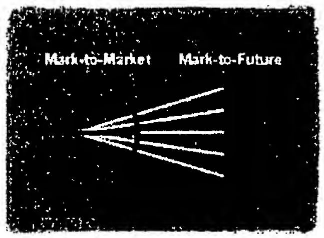

Choose a wide range of scenarios to describe all possible future events. Scenarios should include extremes and those that contradict popular opinion. Then, assign a likelihood to each scenario.

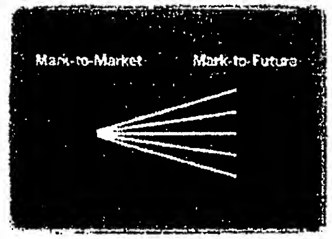

Value your aged portfolio at the horizon under each future scenario. Once the values have been computed, it is possible to determine the Regret and Upside. The **Risk-Adjusted Value** of the portfolio is simply the Upside—Regret.

composition has been determined. This makes computation of marginal near-time risk/reward measurement feasible, even for large institutions. With current methods this is impossible, since prior to a risk calculation one needs to assemble all portfolios centrally, which is extremely time consuming and fraught with difficulty.

Most importantly, Mark-to-Future will break down significant barriers to the risk service bureau business and make risk management outsourcing a reality, by making it possible for institutions to obtain risk management services without divulging their portfolios. This will level the playing field and put smaller institutions on an equal footing with even the biggest of banks. The same risk management analytics, formerly very expensive and available to only the largest financial institutions, will be easily accessible. Ultimately, this will lead to far more extensive and better risk management worldwide.

Mark-to-Future will also revolutionize application development. Today a risk management system must be all things to all people. Mark-to-Future eliminates the need for monolithic applications; instead, software can be created to meet the specific needs of each and every business unit within financial institutions. These 'thin client' applications will be network-enabled and will access pre-computed Mark-to-Future data. Because applications can be built quickly and easily and do not require a large investment, they may be modified quickly to adapt to an ever-changing risk management landscape.

This document is the result of a collaborative undertaking by an extremely knowledgeable, dedicated and talented workforce at Algorithmics over a 10-year period. It is not abstract theory. The ideas presented here have been tried and tested and are available in our products today.

The Mark-to-Future framework is the latest milestone in an incredible journey to solve the future of risk/reward. This document is the first in a series of documents that will describe the Mark-to-Future methodology in detail. We anticipate that this framework will evolve and grow over time and consequently welcome and look forward to your feedback. Updates to the methodology will be posted on www.mark-to-future.com.

Ron S. Dembo, Toronto, November 29, 1999

Part 1
Overview

We provide a description of the framework advocated by U.S. regulators. This, in conjunction with a brief history of risk management systems and regulatory and supervisory requirements, serves as an introduction for the Mark-to-Future framework. A detailed presentation and discussion of the benefits of the framework are provided. It concludes with a brief overview of the methodology.

1

A Call for a FRAMEWORK

Governor Meyer, of the US Federal Reserve Board (1999), has vigorously advocated increasing both the scale and the scope of public disclosure of information for various risk categories, most notably credit exposure and credit concentration. However, there are many challenges when coupling minimum capital regulation to the internal risk-profiling system of a bank. While he has identified internal bank systems as the first lines of defense in the prevention of undue risk-taking, he is quick to point out that supervision and regulation should not duplicate the efforts of increasingly sophisticated, internal best-practice risk management systems.

Still, supervisors and regulators must have confidence that these risk management systems are sufficient. In a recent address, Alan Greenspan, Chairman of the US Federal Reserve Board (Greenspan 1999a), questions the suitability of formulaic approaches and espouses the development of a framework that is adaptable to the inevitable changes in supervision and regulation:

We are striving for a framework whose underlying goals and broad strategies can remain relatively fixed, but within which changes in application can be made as both bankers and supervisors learn more, as banking practices change, and as individual banks grow and change their operations and risk-control techniques. Operationally, this means that we should not view innovations in supervision and regulation as one-off events. Rather, the underlying framework needs to be flexible and to embody a workable process by which modest improvements in supervision and regulation at the outset can be adjusted and further enhanced over time as experience and operational feasibility dictate. In particular, we should avoid mechanical or formulaic approaches that, whether intentionally or not, effectively 'lock' us into particular technologies long after they become outmoded. We should be planning for the long pull, not developing near-term quick fixes. It is the framework that we must get right. The application might initially be bare-boned but over time become more sophisticated.

Mark-to-Future (MtF), a proposed new risk-reward management framework, systematically satisfies the requirements set out by both Meyer and Greenspan.

MtF is a comprehensive, yet flexible, methodology that links disparate sources of risk and provides a means for calculating the risk/reward trade-off for all types of risk within a single, unified framework. In a MtF framework, linking these disparate sources of risk is natural and occurs at the scenario level. Scenarios that embody simultaneous changes in market, credit, and liquidity states naturally provide correlated, consistent output, from which risk measures that link these sources of risk can be calculated. New sources and measures of risk, therefore, can be easily incorporated. Further, innovations in supervision and risk management best-practice can be readily accommodated: neither practitioners nor regulators are locked into so-called formulaic approaches.

As Greenspan contends, risk management is all about the fundamentals: it is the framework that we must get right. In the remainder of this document, we introduce this new MtF methodology in great detail and describe how this framework can significantly advance risk management practice and oversight.

The Risk Management Landscape

The Evolution of Financial Markets

Over the last two decades, the impact of telecommunications and information technology on capital markets has been staggering. The complexity of financial instruments has risen dramatically: the more individuals trade, the more innovative financial securities become. This globalization of financial markets and the securitization of non-traded assets has dramatically expanded the range of investment opportunities. The Internet is igniting a global transformation, where Web-based finance will become the dominant force in the new world economy.

Financial innovation has also increased the number of investment choices. As the space of attainable risk-reward combinations has widened, markets have become more complete and efficient. However, the potential for suffering significant losses has also grown. It is relatively easy for a handful of traders to create very large, leveraged positions, as the near-collapse of Long Term Capital Management and the demise of Barings demonstrated. Financial institutions and their clients are often unaware of the magnitude of potential downside losses.

The relevance of modern risk management techniques to financial institutions and their clients has increased in parallel. Risk measurement techniques that combine the risks of different asset classes and business units consistently at the enterprise level are now considered essential.

The Evolution of Risk Measurement and Management

In response to these market trends, the financial industry has developed a comprehensive set of best-practice guidelines for the measurement and management of market risk at an enterprise level. The twenty recommendations for dealers and end-users of derivatives published by the Group of Thirty (G-30) in 1993 (G-30 1993) still serve as a reference point for all subsequent regulatory and industry efforts to refine and advance best-practice. The RiskMetrics methodology (JP Morgan 1996) and quantitative regulatory standards (BIS 1996, 1997) have become important benchmarks for the estimation of market risk.

Deficiencies in credit risk management have caused much of the recent turmoil in the financial markets. In many instances, such as in the recent Asian crisis, a market risk provoked significant credit and liquidity events. Market-, credit- and liquidity-sensitive exposures pose special challenges to the modeling of counterparty credit exposures when exposures are contingent on unknown future events.

The Counterparty Risk Management Policy Group (CRMPG) has recently published a report on advancing counterparty risk management practices (CRMPG 1999). Both the BIS and the CRMPG stress the importance of modeling contingent exposures correctly (BIS 1999c, CRMPG 1999). In response to these challenges, the estimation of counterparty credit exposures is moving from a static add-on approach (BIS 1988) to an extended add-on framework (ISDA 1998), and towards a statistical estimation of potential future exposures (BIS 1999c).

Common industry standards for the modeling of portfolio credit risk continue to evolve. Several methodologies such as CreditMetrics (1997), CreditRisk+ (Credit Suisse Financial Products 1997) or CreditPortfolioViewer (Wilson 1997a, 1997b) have been proposed. The problem with all of these methods is that they treat credit risk in isolation. They ignore the fact that credit and market risk are inextricably linked. Treating these risks independently of one another makes the proper modeling of credit derivatives and market risk exposures next to impossible.

Following the near-collapse of Long Term Capital Management in 1998, industry workgroups and regulators have promoted the proper measurement of liquidity risk (CRMPG 1999, BIS 1999a, BIS 1999b). The CRMPG distinguishes two categories of liquidity risk: (1) asset liquidity, the ability to sell or unwind positions; and (2) funding liquidity, the ability to meet obligations when due (CRMPG 1999).

In addition, the BIS has recently proposed principles for the management of credit risk and for the integrated management of market, credit and liquidity risk (BIS 1999d, 1999e).

In summary, then, it is clear that regulatory requirements and industry best-practice are converging on the need for an integrated risk measurement framework that

- aggregates risk consistently across many asset classes
- captures the interaction of market, credit and liquidity risk
- combines statistical risk measures for market and credit risk and includes stress testing
- supports risk oversight for the enterprise and for local trading desks. ◎

A Brave New World of Risk

A FRAMEWORK FOR MEASURING RISK
(not a risk measure)

FORWARD-LOOKING
(takes account of future events, scenarios)

PROPER FORWARD VALUATION
(scenario dependence, forward P&L)

RISK OVER TIME

DYNAMIC PORTFOLIOS

NATURALLY LINKS MARKET, CREDIT, LIQUIDITY, LEGAL AND OPERATIONAL RISK

COMPUTES RISK / RETURN TRADE-OFF

FOCUS IS ON INSTRUMENTS, NOT PORTFOLIOS

As we have seen, nowhere is that more true than in financial markets and in the risk management industry. The need for a risk management framework that is adaptable to these on-going and inevitable changes has never been greater.

Enter Mark-to-Future. Mark to Future is the first truly forward-looking framework for risk/reward measurement and management. Unlike covariance-based approaches which assume a portfolio remains constant over time and that returns or values are normally distributed Mark-to-Future makes no such assumptions. The framework accounts for all events between the present and some future time horizon, including the effects of settlement, cash accounting and future value of aging.

To look forward in time, one needs to consider the portfolio under many different scenarios (Figure 3.1). However, many instruments have a valuation that is scenario-dependent. Consider an AA fixed-rate corporate bond. The price of the bond is sensitive to market and credit risk. Scenarios include simulated changes in market rates, credit spreads and credit downgrades. When a credit downgrade to say, A occurs, the simulated A credit spread curve rather than the AA curve has to be used together with the market rate in that scenario. If the same AA curve were used under all scenarios, or if the market or spread volatility were not considered, the true risk of the instrument would not be captured and the end result would be misleading risk information. To properly

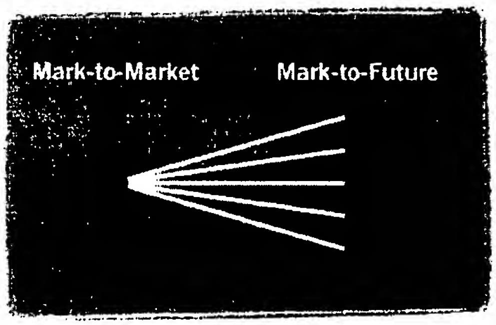

Figure 3.1: The MtF framework considers the portfolio under many different scenarios at future points in time

value the risk of such instruments, the reinvestment risk of the settled coupons must also be taken into account.

Mark-to-Future is not a static framework with a view towards a single horizon or a model for extrapolating risk. It explicitly incorporates the passage of time (Figure 3.2). Consider a 10-day Zero-Coupon bond. The 10-day VaR should be 0; that is, there is no risk at 10 days. However, a methodology that computes forward risk by simply extrapolating risk forward in time, would generate a positive VaR number and result in an incorrect allocation of capital. With this new paradigm, trading strategies such as dynamic hedging can be modeled through event-based rules. The end-result is a portfolio whose composition is both time and scenario dependent. Thus, the risk of trading desks whose positions change over time and under scenarios can be measured.

Scenarios, not summary statistics, serve as the core input information in a Mark-to-Future framework. Except in extreme cases, very few practitioners can determine whether a summary statistic—such as a covariance matrix—is reasonable or not. Further, only for a very specific set of scenarios can a single parameter truly describe the robustness of that distribution. In risk methodologies that use a summary statistic as the key information input, a set of scenarios is implied. The Mark-to-Future framework, in contrast, forces one to be explicit about the scenario choice. These scenarios must span all possible future outcomes over relevant time steps and be forward-looking in nature.

History is an excellent source for reasonable event sequences. These events can be considered reasonable because they have actually occurred. If there are no fundamental shifts in the underlying market structure and there is a long history, historical scenarios span the range of future events. To be truly forward-looking, though, these scenarios should contain events not currently in the information set. Yet, some markets, such as those for Internet stocks or emerging markets, are too new to have sufficient history.

With some careful analysis, however, and a broad perspective on where information for scenario generation can be found, it is possible to find ways to supplement the 'normal' scenarios that might be obtained from recent history. By explicitly selecting scenarios, it is possible to incorporate both normal and stress scenarios and de-couple scenario choice from the choice of risk measures.

One of the more challenging aspects of risk management, as observed earlier, is linking market, credit and liquidity risk. In a Mark-to-Future framework, linking these disparate sources of risk is natural and occurs at the scenario level. A scenario, for example, might include credit downgrades and interest rate changes. Once that scenario is defined, and a model for valuation agreed upon, the Mark-to-Future methodology links these disparate risks. It is then possible to generate consistent Mark-to-Future output, from which risk measures that link these sources of risk can be calculated. The framework simplifies the approach needed to integrate these risks/rewards for capital allocation purposes.

It is clear, then, that the Mark-to-Future methodology unequivocally 'answers the call' for a new risk management framework (Appendix A). It incorporates forward-looking scenarios that account completely for all potential future events and their effect on the future value of financial instruments. Since the framework incorporates risk over time, the change in portfolios can be measured. Most importantly, Mark-to-Future makes it possible to link disparate sources of risk within a single, unified framework.®

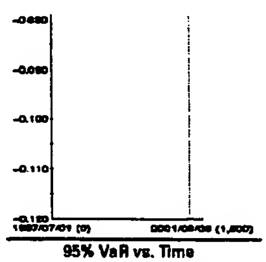

Figure 3.2: MtF VaR of Bond Call with Reinvestment: This graph highlights the importance of computing risk over time. The largest exposure may occur before the horizon, shown here as the dotted line

THE PRODUCTION OF DAILY CONSOLIDATED, CONSISTENT, AND COMPREHENSIVE ENTERPRISE RISK MANAGEMENT REPORTS HAS CHALLENGED MANY SOPHISTICATED FINANCIAL INSTITUTIONS. THE MARK-TO-FUTURE PARADIGM ADDRESSES NOT ONLY THE METHODOLOGICAL ISSUES, AS DEMONSTRATED IN THE PREVIOUS CHAPTER, BUT ALSO THE TECHNICAL AND IMPLEMENTATION ISSUES OF A GLOBAL RISK MANAGEMENT FRAMEWORK

At the heart of the MtF framework is the generation of a 3-dimensional cube of MtF values (Figure 4.1). This cube contains all the information required to compute risk/reward. MtF separates the simulation (cube computation) from the statistics used to compute risk and reward measures. The same MtF Cube may be used to compute multiple scenarios.

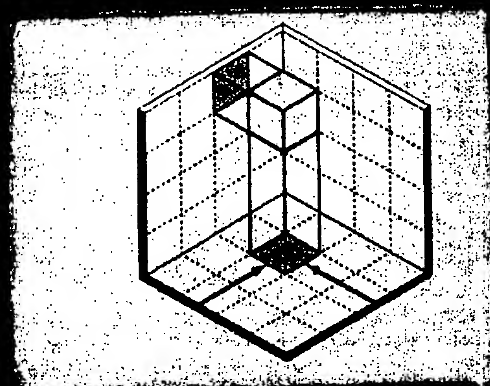

Figure 4.1. Basis Mark-to-Future Cube

A key benefit of the MtF framework from an architectural perspective is that when scenarios and time steps are consistent between MtF cubes, the MtF values are strictly additive along the instrument dimension. Thus, in a MtF world, the focus is on instruments, not portfolios.

The Bottom

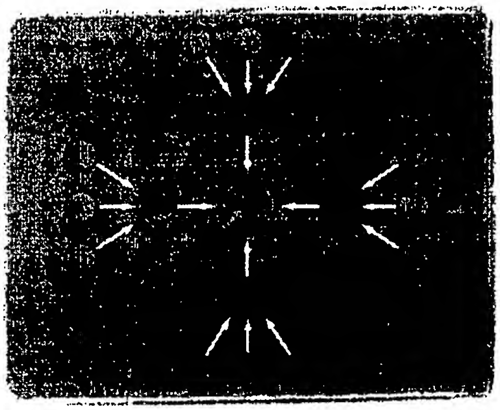

Figure 4.2: Distribution of MtF generation over multiple servers

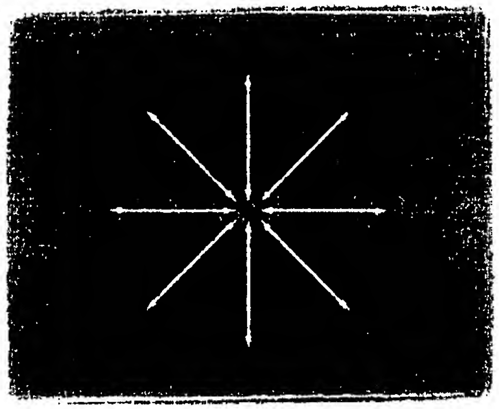

Figure 4.3: Consolidation of end-of-day positions in a centralized system

By calculating the risk of the underlying instruments, massive reductions in computation time are possible, since the risk of the portfolio is a simple linear combination of the risk of each instrument within the portfolio. This allows for a scalable solution to be achieved with a decentralized architecture and process.

Further, the creation of a global MtF Cube can be accomplished within a distributed environment and computed at different points in time through the generation of multiple component MtF Cubes that are then aggregated. The production of these cubes can be distributed over multiple servers; these servers may be centrally located or geographically dispersed (Figure 4.2). These cubes appended together form a consolidated basis. This radically alters the potential architecture for a global risk management implementation.

No Need to Consolidate End-of-Day Positions

Traditional, centralized risk management requires that all end-of-day positions be confirmed and assembled before risk can be calculated (Figure 4.3). Thus, the speed with which you can process an enterprise risk management (ERM) analysis is inherently limited to the speed with which you can gather all of the data on portfolios.

Typically, it might take 6 hours or more to assemble positions in a warehouse, assuming nothing has gone awry. Only then does risk computation begin. Thus, even if the risk calculation time was negligible, a run would take over 6 hours. This precludes any type of true marginal risk computation in real time.

In a MtF framework, the enterprise-wide process for the consolidation of risk and reward measurement need not be delayed by the reconciliation of portfolio positions across the organization. The generation of the Basis MtF Cube can begin as soon as market price data becomes available.

Thus, the majority of the simulation work may be completed before positions are known. This can radically reduce the overall computation time for an enterprise risk calculation.

Line

Compute Once, Use Many Times

Many portfolios have common instruments. Simply consider the number of times throughout the day that the same bond is evaluated for many different purposes in an investment bank.

By focusing on instrument valuation by scenario rather than on portfolios, MtF takes advantage of the fact that if the scenarios and time horizons are consistent, the MtF vectors are additive; that is, risk is linear in MtF space. Any combination of the MtF values of the individual portfolios has a MtF that is a sum of their individual MtF values (Figure 4.4). By storing the MtF of a comprehensive set of basis instruments, the value of any portfolio is simply a linear combination of the value of these instruments, regardless of whether they are linear or non-linear.

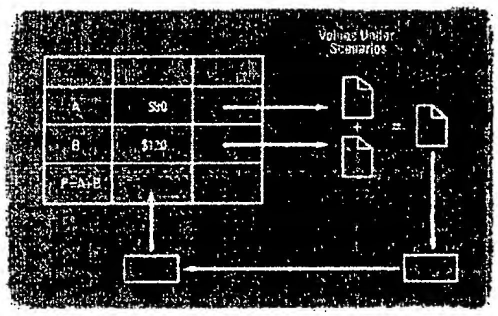

Figure 4.4: The MtF solution: Values under scenarios and through time are additive (i.e. MtF tables).

The MtF of a basis instrument, therefore, only needs to be computed once. MtF enables a 'compute once, use many times' approach which radically reduces the computing power required for global and local risk reporting and intra-day analysis.

Aggregating Along Multiple Portfolio Hierarchies

Typically, a global risk report includes summary risk information aggregated along multiple portfolio lines—by business unit, by product and risk factor type or by counterparty. The portfolio MtF is simply the sum of the instrument MtF values, weighted by the size of the position, along each of the stipulated portfolio lines.

Accordingly, the MtF of any alternate portfolio hierarchy is the simply a weighted sum of any set of instrument MtF values. Thus, 'on-the-fly' drill-down and aggregation analyses can be performed in the course of supporting daily activities.

Correction of Errors and Omissions

Inevitably, omissions, updates and errors in preliminary positions occur. The draft portfolio is a portfolio with MtF values that exclude the missing position. Although processing occurs before confirmation of end-of-day positions in a MtF framework, omissions are easily corrected. The correct MtF values of the portfolio can be calculated by executing an action to add the MtF vector of a missing, or reconciling, position of the correct size to the MtF values of the draft portfolio.

Consider a position in a bond future for which the wrong position size has been entered. The preliminary position is a portfolio with a MtF value that includes the wrong position. The correct portfolio MtF values can be calculated by executing a reconciliation action to reverse the MtF values of the incorrect position, rescale the product MtF cube by the right position size, and add the resulting MtF values to the portfolio MtF cube. Because MtF values are additive, a full revaluation is avoided.

As a second example, consider a position in a bond for which the wrong coupon rate has been entered. Again, the correct portfolio MtF value can be calculated by executing a reconciliation action to reverse the MtF values for the position in the bond with the wrong rate and simultaneously add the MtF values for the position of the same size in a bond with the correct rate.

In these two cases, the reconciliation is a simple matter of adding a new MtF vector, or reversing one vector and adding another. More complicated reconciliation actions are composed of these simple components. For example, if the coupon rate for the bond underlying a bond option is in error, a series of reconciliation actions would

- calculate the instrument MtF for the correct bond
- calculate the instrument MtF for the bond future
- calculate the instrument MtF for the bond option
- reverse the MtF vector on the incorrect bond option
- add a new MtF vector for a bond option with the correct rate.

Thus, any errors, omissions and updates in positions can be corrected as a subsequent adjustment exercise.

Marginal Deal Analysis

The additivity property of the MtF framework also enables near-time marginal risk calculations at the desk level. By precomputing the MtF data and releasing it across a global network, consistent risk numbers are easily distributed across the organization.

This has significant business benefits. Consider the incremental impact of a new deal on the credit exposure to two alternative counterparties. While the potential exposure of the new deal is the same on an absolute basis for each counterparty, the deal may act as a natural hedge for the exposure to one counterparty

on a marginal basis while increasing the exposure to the other counterparty.

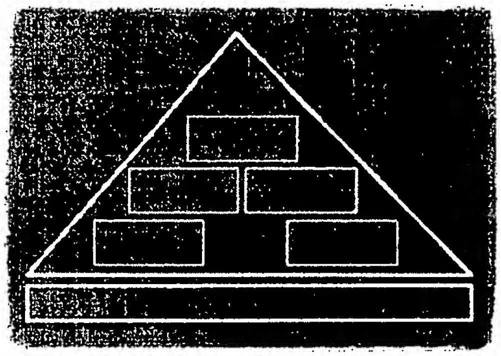

Figure 4.5: The 'Disposable' Application: The MtF paradigm makes it possible to move away from the traditional paradigm where a complex and monolithic application must be all things to all people

Lightweight Applications

With this new approach, it is now possible to create applications designed to meet the needs of individual business units within a financial institution (Figure 4.5). Applications—such as a simple VaR calculator for a bond desk or a mark-to-market calculator for a mortgage-backed securities desk—are much quicker and simpler to create when MtF data is available. These applications can be built within a short period of time and revised and adopted quickly to meet changing business needs. Because of the ease with which these can be built, customization to the specific needs of individuals within the institution can be met.

Security Issues Surrounding Proprietary Holdings

The current portfolio is a necessary input in traditional risk measurement; all analytics are applied to the portfolio object. In a traditional 'pull' system, service bureaus offering risk reports to their clients must collect the details of proprietary holdings to prepare these reports.

A financial institution or service bureau that has adopted MtF can distribute an adequate subset of the Product MtF Cube to business units or clients. Using custom applications, the business units or clients can then calculate the necessary risk measures at the desired level of aggregation. They need not transmit sensitive information concerning their portfolio holdings. This is likely to have a profound effect on the growth of risk outsourcing and service bureaus. It will also level the playing field. Small institutions will have access to the same quality risk tools as large banks.

Dynamic Strategies

Modeling Dynamic Risk Over Long Time Horizons

The risk of a portfolio can vary greatly over time. To measure risk over time, it is important to model the change in portfolio composition. These changes may be due to portfolio aging, which leads to settlement and re-investment, with associated funding and collateral costs. Portfolio changes may also be a result of dynamic strategies designed to change the characteristics of a portfolio in response to market events.

It is also the case that for longer horizons that are appropriate for pension funds, asset managers and strategic decisions, the dynamics of a portfolio cannot be ignored. To measure risk dynamically, simulation over multiple periods is necessary (Figure 4.6).

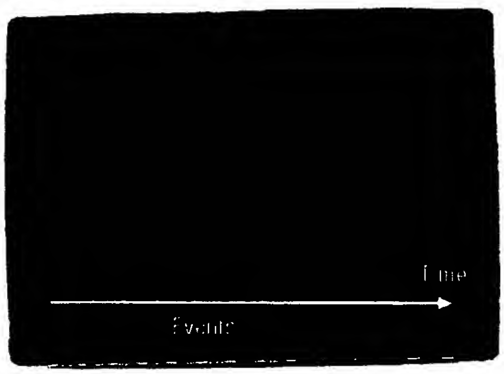

Figure 4.6: MtF tracks risk over time (Funding, Reinvestment, Liquidity, Hedging, etc.)

To model the dynamics of a portfolio, the use of a regime or dynamic portfolio strategy is necessary. A regime is a set of rules, which describe the conditions under which the composition of a portfolio is to be changed at a specified time. Delta hedging is an example of a regime that is based on recent history and varies depending upon the scenario. The rule might be

If the net delta of the portfolio and the hedge is more than 3% away from the desired net delta of zero target, rebalance the hedge using futures.

A more complicated rule, might have an additional clause such as

If margin on the derivatives book is below s, liquidate sufficient bonds to raise it to S.

Any logical expression based on information known deterministically at the point at which the expression is to be evaluated, is valid.

This way of modeling the dynamics of a portfolio has significant computational advantages.

Figure 4.7: A MtF regime

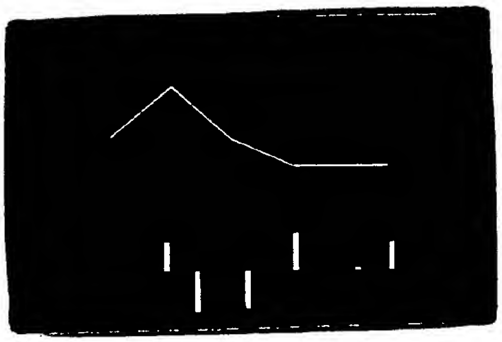

Figure 4.8: MtF takes all portfolio effects into account, including any event that might affect the forward value. Proper cash and collateral accounting is required for MtF.

It presumes that the number of possible interesting regimes is small and then enumerates over all of them. For a trading desk or an asset/liability committee, it is reasonable to assume that there are only a small number of interesting regimes. By making the regimes a deterministic function of the scenarios, what is inherently an intractable multi-stage, multi-period stochastic optimization problem with thousands of risk factors, is reduced to an easily solved single stage, multi-period problem (Figure 4.7).

Portfolio Aging, Settlement and Reinvestment

As a portfolio ages, it spins off cash which is settled or re-invested into other instruments. The total risk in a portfolio includes re-investment risk. Dynamic portfolio positions can be adjusted to credit and debit cash and collateral (or physical) accounts as positions mature. Re-investment risk can be measured by accounting for the changes to the cash and collateral accounts, by scenario. A necessary aspect of MtF is, therefore, the proper forward accounting of cash and collateral required to support these positions (Figure 4.8).

Measuring Liquidity Risk

Liquidity risk can be measured by creating scenarios on appropriate risk factors and by defining the appropriate liquidation regimes. A regime for liquidation describes parameters for determining the speed and cost with which a portfolio can be liquidated. Liquidation may also depend on simulated liquidity risk factors such as trading volume. Each scenario path generates MtF values and a fluctuating cash/collateral account. At any time point, therefore, there is a distribution of MtF values and of the cash/collateral required to support this portfolio. The distribution of MtF values provides a measure of market risk adjusted for asset liquidity. The distribution of the changes to the collateral/cash account over time is a measure of funding liquidity risk (Figure 4.9).

The graph depicts the risk of a portfolio containing US treasuries, Gilts, liquid stocks, illiquid stocks and illiquid corporate bonds (Figure 4.10). The view to the left illustrates the value of the cash account over time and by scenario. The histograms to the right aggregate these values into a frequency distribution at three different

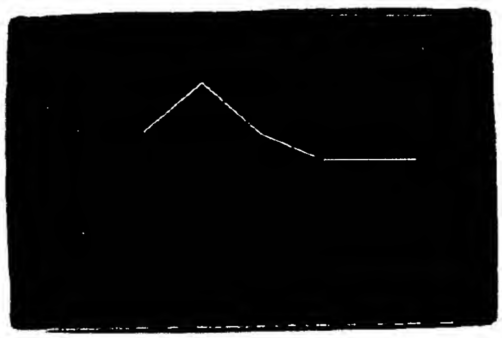

Figure 4.9: Funding Liquidity Risk: Each scenario path creates a fluctuating cash/collateral account

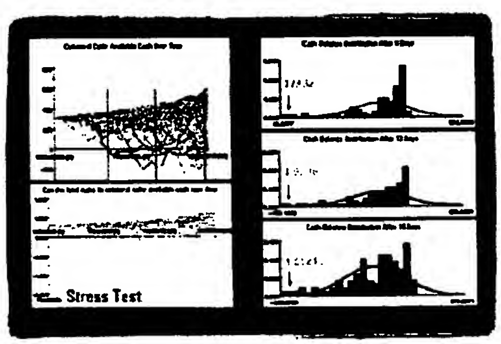

Figure 4.10: Collateral Dynamic Risk in the MtF Framework. The distribution of changes to the collateral/cash account is a measure of funding liquidity risk from which time dependent 'liquidity VaR' can be computed.

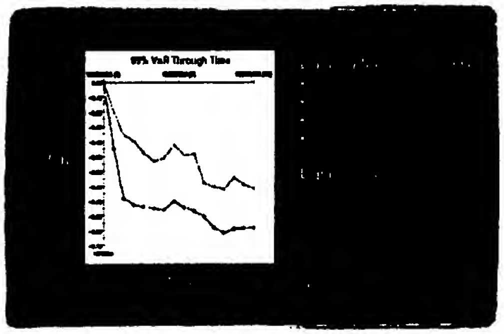

Figure 4.11: Various liquidation regimes affect the VaR over time.

time steps. To illustrate liquidity effects, we demonstrate how various liquidation regimes might affect the value of VaR over time (Figure 4.11).

Integrated Market, Credit and Liquidity Risk

One of the more challenging aspects of risk management is linking market, credit and liquidity risk. In a MtF framework, linking these disparate sources of risk is natural and occurs at the scenario level. Scenarios that embody simultaneous changes in market, credit and liquidity states naturally provide correlated, consistent MtF results, from which risk measures that link these can be calculated. A scenario, for example, might include both credit downgrades and interest rate changes. Once that scenario is defined, and a pricing model specified, MtF links these risks and accounts for the inherent correlation.

Traditional risk management systems often separate stress testing and standard market risk calculations. In a MtF framework, stress tests and normal market events are contained within the same MtF Cube. The same calculation steps apply to both categories of scenarios consistently. Only when MtF values are post-processed to calculate risk measures are the appropriate probability weights applied to scenarios in each category. The MtF cube, without choice of probability weights, contains very useful information for risk management purposes.

Any computation of risk/return trade-offs and risk-adjusted return is based on the same pre-computed MtF cube, whether simple market risk scenarios, stress tests or simultaneous market and credit risk scenarios are specified. In this sense, MtF is a unifying framework that simplifies the integration of risk and reward measurement for capital allocation purposes.

Consistent Measures of Risk and Reward

The MtF framework links disparate sources of risk, and moreover, provides an integrated framework for analyzing the trade-off between these risks and reward.

Consider simulating the value of a portfolio over a single period. The MtF simulation is performed on a set of basis instruments, over a set of scenarios and an appropriate time horizon. The financial products are mapped to this set of instruments and the portfolio is mapped to the financial products. The result is a set of MtF values for the portfolio, one for each scenario.

Some of these scenarios will have resulted in gains and some in losses. The gains and losses could be absolute or computed with respect to a benchmark, the benchmark portfolio itself being subject to the same simulation. The gains and losses provide all the information required to compute a forward-looking measure of risk or reward (Figure 4.12).

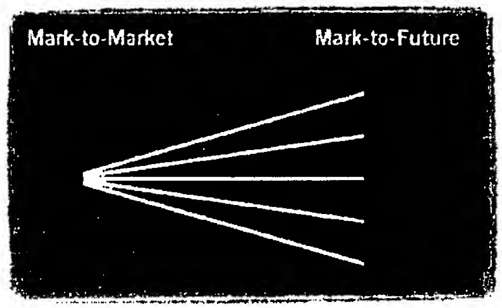

Figure 4.12: By calculating the Upside and the Downside, it is possible to compute a forward-looking measure of risk-adjusted return

If probabilities are associated with each scenario the result is a distribution of unrealized gains and losses. Many risk or reward measures can be derived from this distribution. For example, the parametric VaR measure is a specified percentile of the downside, using the mark-to-market value as the benchmark. It is also easy to evaluate multiple choices for the probability weights, since the MtF Cube remains invariant.

The Put/Call Efficient Frontier (Dembo 1999) is an important new means of measuring both risk and reward. The determination of the Put/Call Efficient Frontier begins with the decomposition of the MtF into its downside losses and upside gains, representing the risk and the reward respectively.

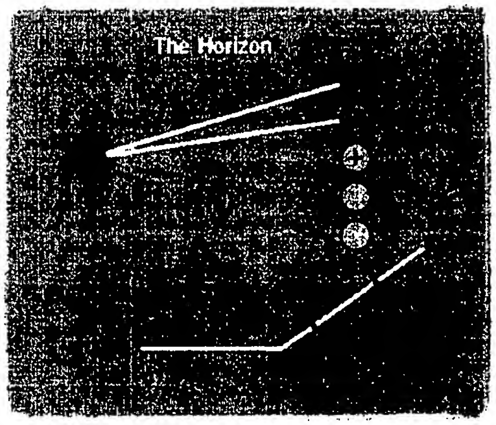

Figure 4.13: The upside has the same payoff as a European call option

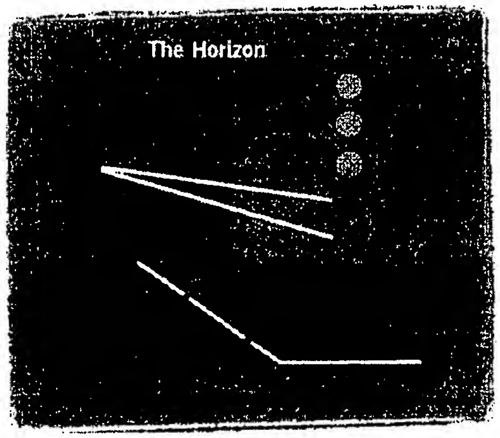

Figure 4.14: The downside has the same payoff as a European put option

The upside has the same payoff as a European call option, with strike equal to the benchmark value and maturity equal to the horizon (Figure 4.13). Similarly, a short position in a European put option, with strike equal to the benchmark value and maturity equal to the horizon, has the same payoff as the downside (Figure 4.14). A value can be placed on the upside and downside of a portfolio by valuing the call and the put. Thus, in a MtF world, any portfolio may be disaggregated into its put and call values.

This is an inherently forward-looking view of risk and reward since the value of a put or call is dependent exclusively on future events, though the past might influence the choice of scenarios.

This leads to a natural way of measuring risk-adjusted performance. The adjusted upside is the value of the call less the value of the put scaled by a risk aversion factor. This risk-aversion parameter, which is assumed to be greater than one, reflects psychological or other non-quantifiable factors. The adjusted upside is a measure of risk-adjusted value as it includes the (perceived) cost of insuring the downside. It forms the basis of how we propose that capital be allocated. The scaled value of the downside is referred to as Regret (Dembo 1998), and is itself a useful risk measure.

The trade-off between risk and return for any portfolio is therefore captured by the trade-off between its put and call value. The Put/Call Efficient Frontier traces the maximum upside for given levels of downside. Equivalently, it may be formulated as the minimum level of downside for fixed levels of upside (Figure 4.15).

The risk and reward of the portfolio today for some given horizon can be calculated by placing a value on this put and call. The methodology for valuing the put and call is presented in Part 2, Chapter 7.

It is important to note that in developing the Put/Call Efficient Frontier no assumptions have been made concerning the sources of risk or about the distributions of securities or payoff functions, as would be the case in a mean-variance framework (Markowitz 1952). If the scenarios contain events that link market, credit and liquidity risk and if the MtF is properly computed, then the Put/Call Frontier will be an efficient frontier that trades these integrated risks against their upside. Correlation between these types of risks is accounted for, since it is implicit in the choice of scenarios.

This further reinforces one of the primary principles behind MtF. By separating the simulation from the risk measure itself, the trade-off between risk and return and the integration of various types of risks may be developed in a generic manner.

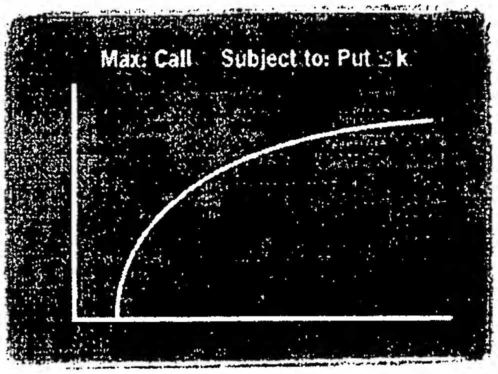

Figure 4.15: Put/Call Efficient Frontier

MtF, then, represents a risk management framework that provides global risk oversight and trading decision support, for end-of-day and intra-day business functions. The advantages derived from this radical new approach are numerous. Time consuming MtF calculations can be staggered and distributed, as resulting values are fully additive. The additivity of the MtF values also facilitates intra-day updates at the trade level. Changes in position units result in a rescaling rather than a recalculation. Once calculated, the MtF values can be re-aggregated into different risk measures or disaggregated for more detailed analysis. Marginal analysis of VaR can be performed, credit exposures and losses can be calculated, all without the need to recalculate MtF values. Custom reports based on different risk measures can be created to meet the specific needs of various users across the enterprise.

The SIX Steps

In the last several chapters, we have provided a detailed description of the MtF framework and have outlined the many advantages adoption of the methodology offers. We now provide a brief introduction to the MtF Cube and describe how it is generated.

The development of the MtF framework involves two critical stages, a computationally intensive **pre-Cube** (or simulation) stage and a less computationally intensive **post-Cube** stage. The pre-Cube stage is associated with the configuration and generation of the Basis MtF Cube while the post-Cube stage is associated with the mapping of portfolios and portfolio regimes onto the basis and the application of post-processing analytics to the resulting portfolio MtF values.

The following six key steps summarize the overall risk-reward assessment process as embodied by the MtF framework. The first three represent the pre-Cube stage and the last three represent the post-Cube (or post-simulation) stage. Each of these steps can be explicitly configured as an independent component of the overall process:

1. Define the basis instruments.
2. Define the scenario paths and time steps.
3. Simulate the instruments over scenarios and time steps to generate a Basis MtF Cube.
4. Map the financial products onto the Basis MtF Cube to form a Product MtF Cube.
5. Map the portfolio regimes onto the Product MtF Cube to form a Portfolio MtF Cube.
6. Post-process to produce risk, reward and performance measures.

In specific applications of the MtF framework, some of these steps may not be present. In particular, when there are few discrete financial products, the mapping of basis instruments to financial products may be 1:1 so that Step 4 would be redundant. In situations where the number of instruments may be large, mapping onto basis instruments helps reduce the dimensionality of the problem. Often, too, it can have a dramatic effect on computational speed, as is the case for large equity or swap portfolios.

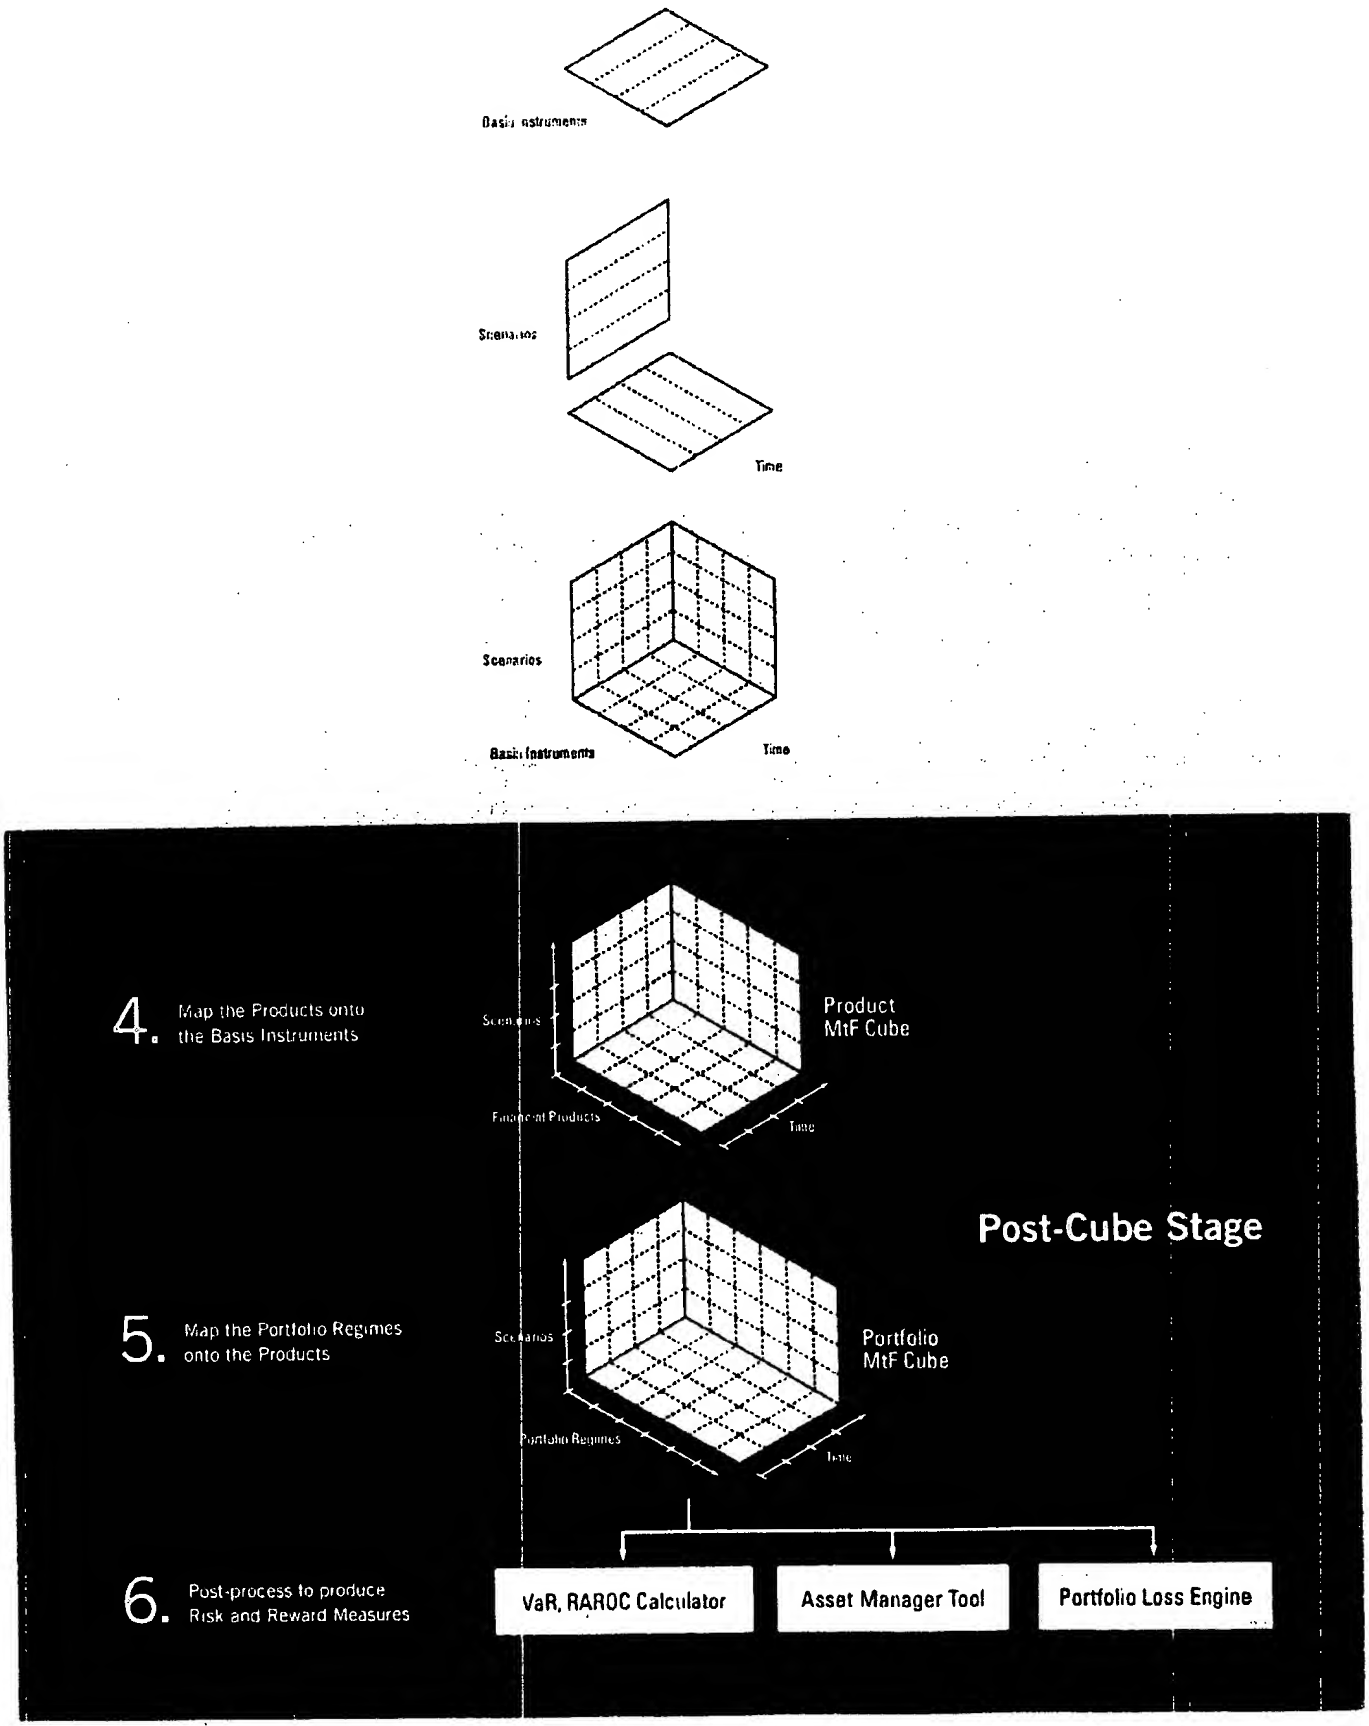
Basis Instruments
Scenarios
Time
Scenarios
Basis Instruments
Time
4. Map the Products onto the Basis Instruments
Scenarios
Financial Products
Time
Product MtF Cube
Post-Cube Stage
5. Map the Portfolio Regimes onto the Products
Scenarios
Portfolio Regimes
Time
Portfolio MtF Cube
6. Post-process to produce Risk and Reward Measures
VaR, RAROC Calculator
Asset Manager Tool
Portfolio Loss Engine

1: Define the Basis Instruments

The first step involves the appropriate selection of instruments required to generate the Basis MtF Cube. Basis instruments may represent actual financial products such as stocks and bonds, synthetic instruments such as a series of zero coupon bonds or, in some cases, abstract indices such as the inputs to a multi-factor model. In most applications, the Basis MtF Cube consists of a mixture of actual products representing exchange-traded securities or over-the-counter (OTC) securities held in inventory and synthetic instruments onto which newly transacted OTC securities may be mapped.

Basis instruments are selected to capture the range of relevant risk factors associated with the valuation of any existing or potential financial product. This choice is largely dependent upon the individual business application and may involve tradeoffs between the quantity of data storage and pricing accuracy. Nonetheless, a judicious choice of basis instruments can mitigate the data storage issue, often with little or no loss in pricing accuracy.

The settlement and reinvestment of cashflows that may occur over a given time horizon must be addressed during the selection of basis instruments. Consequently, basis instruments must be flexible enough to represent the 'total return' of a financial product by capturing any cash flow settlement and by incorporating a choice of reinvestment possibilities.

2: Define the Scenario Paths and Time Steps

Once the basis instruments have been defined, the scenarios and associated time steps must be selected to form the other two dimensions of the Basis MtF Cube. Scenarios describe the evolution of possible risk factor levels, including market, credit and liquidity risk factors in the future.

3: Simulate to Generate a Basis MtF Cube

Once the basis instruments have been selected and the scenarios and time steps defined, the Basis MtF Cube is generated by applying a pricing model that produces a mark-to-market valuation for each scenario. The generation of the 'Scenario x Instrument x Time Step' Basis MtF Cube in these first three steps completes the computationally intensive pre-Cube stage of the MtF framework. The MtF values contained in the Basis MtF Cube serve as the raw data for the next stage of processing, which involves the mapping of the actual financial products if necessary, the portfolio holdings and the measurement of risk and reward. The next three steps describe the post-Cube stage of the MtF framework.

4: Map the Financial Products onto the Basis MtF Cube to Form a Product MtF Cube

Next, a mapping exercise is required to produce MtF values for the virtually unlimited number of financial products that exist. This exercise calculates a MtF value for a given financial product as a function of the MtF values of the basis instruments contained in the Basis MtF Cube.

The result of this mapping exercise is the creation of a **Product MtF Cube** containing the MtF values of all financial products of interest. Unlike the Basis MtF Cube, the Product MtF Cube need not exist in a physical sense as it is completely defined by the product mappings themselves. That is, the cube need not be saved, it only needs to be calculated as required.

In many cases, the product mappings produce perfect MtF value replications for the financial products across all scenarios and time steps. In certain cases, however, the mapping of financial products onto basis instruments will result in imperfect replication. This may be because a less complex mapping function has been chosen for a specific application. It may also occur because the Basis MtF Cube does not contain basis instruments associated with all relevant risks that affect the value of an instrument. In the latter case, the mapping must include basis instruments as well as those risks that are not contained within the Basis MtF Cube.

5: Map the Portfolio Regimes onto the Product MtF Cube to Form a Portfolio MtF Cube

To determine the MtF values of actual portfolio regimes, the fifth step in the MtF framework involves the mapping of portfolio positions onto the Product MtF Cube. The result of this second mapping exercise is the creation of a Portfolio MtF Cube containing the MtF values of all portfolio regimes of interest. Similar to the Product MtF Cube, the **Portfolio MtF Cube** need not exist in a physical sense as it is completely defined by the combined portfolio and product mappings.

Since the MtF framework tracks risk over time, the mapping of portfolio positions onto financial products need not be static. Significantly, portfolios may be dynamically mapped onto products on the basis of pre-determined or conditional regimes. Regimes are also known as dynamic portfolio strategies. Under a **Predetermined Regime** the portfolio positions are independent of the contents of the Basis (or Product) MtF Cube. Within this category of portfolio strategies, the quantity of holdings to be mapped onto each financial product at each time step is defined by a simple position schedule.

Under a **Conditional Regime** portfolio positions are a function of the contents of the Basis (or Product) MtF Cube. Within this category of portfolio strategies, the quantity of holdings to be mapped onto each financial product at each time step is determined as a function of the MtF values under given scenarios and time steps. Under a conditional regime, such portfolio strategies as delta hedging, immunization, reinvestment, collateral and rollover strategies may be assessed in terms of their individual risk and reward profiles.

6: Post-Process the Portfolio MtF Cube to Produce Risk and Reward Measures

The Product MtF Cube contains MtF values for portfolio regimes across all scenarios for each time step. The resulting distribution of portfolio outcomes provides a robust framework for the assessment of risk and reward. The actual risk and reward measures invoked are calculated strictly as a post-processing exercise. In contrast to traditional methodologies, the preferred risk-reward measures calculated in the post-Cube stage are completely de-coupled from the scenario assumptions made in the pre-Cube stage.

In the post-processing step, the desired risk-reward measures as well as other quantitative analytics that may be applied to the Portfolio MtF Cube are calculated. These post-processing analytics will typically be embedded in a series of 'lightweight' task-oriented user applications.

This concludes the high-level overview of the six steps of the MtF framework. We now turn to a detailed, step-by-step discussion of the methodology.

Part 1: References

Bank for International Settlements, *International Convergence of Capital Measurement and Capital Standards*, Basle Committee on Banking Supervision, July 1988. Available at http://www.bis.org/publ/

Bank for International Settlements, *Amendment to the Capital Accord to Incorporate Market Risks*, Basle Committee on Banking Supervision, January 1996. Available at http://www.bis.org/publ/

Bank for International Settlements, *Explanatory Note: Modification of the Basle Capital Accord of July 1988*, as amended in January 1996, Basle Committee on Banking Supervision, September 1997. Available at http://www.bis.org/publ/

Bank for International Settlements, 1999a, *Banks' Interactions with Highly Leveraged Institutions*, Basle Committee on Banking Supervision, January 1999. Available at http://www.bis.org/publ/

Bank for International Settlements, 1999b, *Sound Practices for Banks' Interactions with Highly Leveraged Institutions*, Basle Committee on Banking Supervision, January 1999. Available at http://www.bis.org/publ/

Bank for International Settlements, 1999c, *Credit Risk Modeling: Current Practices and Applications*, Basle Committee on Banking Supervision, April 1999. Available at http://www.bis.org/publ/

Bank for International Settlements, 1999d, *A New Capital Adequacy Framework*, Basle Committee on Banking Supervision, June 1999. Available at http://www.bis.org/publ/

Bank for International Settlements, 1999e, *Principles for the Management of Credit Risk*, Basle Committee on Banking Supervision, July 1999. Available at http://www.bis.org/publ/

Counterparty Risk Management Policy Group (CRMPG), 1999, "Improving counterparty risk management practices," June 21, 1999. Main report, appendices and glossary available at http://www.crmpolicygroup.org/

CreditMetrics: The Benchmark for Understanding Credit Risk, Technical Document, 1997, New York, NY: JP Morgan Inc. Available at http://www.riskmetrics.com/cm/index.cgi

Credite Suisse Financial Products, 1997, *CreditRisk+: A Credit Risk Management Framework*, New York, NY.

Dembo, R.S., 1998, "Mark-to-Future: A Consistent Firm-Wide Paradigm For Measuring Risk and Return," Risk Management and Analysis, Vol. 1: Measuring and Modeling Financial Risk, Carol Alexander (ed.), New York: John Wiley & Sons Ltd.

Dembo, R. S., 1999, "The Put/Call Efficient Frontier," Working Paper, Algorithmics Incorporated.

Dembo, R. S., 1991, "Scenario Optimization," *Annals of Operations Research*, 30: 63-80. Also appeared as US Patent Number: 5148365, dated September 15, 1992 (filed August 15, 1989).

Dembo, R. S., "Method and apparatus for optimal portfolio replication," US Patent Number: 5799287, dated August 25, 1998 (filed May 24, 1994).

Greenspan, A., 1999, "The evolution of bank supervision," speech to American Bankers Association, Phoenix, Arizona, October 11, 1999. Available at http://www.bog.frb.fed.us/boarddocs/speeches/1999/

Group of Thirty (G-30), 1993, *Derivatives: Practices and Principles*, Global Derivatives Study Group, Washington, D.C. (http://www.group30.org)

Markowitz, H.M., 1952, "Portfolio selection," *The Journal of Finance*, 7(1): 77-91.

Meyer, L., 1999, "Implications of recent global financial crises for bank supervision and regulation," speech before the International Finance Conference, Federal Reserve Bank of Chicago, Chicago, Ill. October 1, 1999. Available at http://www.bog.frb.fed.us/boarddocs/speeches/1999/

Wilson, T., 1997a, "Portfolio credit risk I," *Risk*, 10(9): 111-117.

Wilson, T., 1997b, "Portfolio credit risk II," *Risk*, 10(10): 56-61.

Zangari, P. and J. Longerstaey, JP Morgan, 1996, *RiskMetrics-Technical Document*, 4th ed., New York, NY: RiskMetrics Group. Available at http://www.riskmetrics.com/rm/index.cgi

Appendix A

In this table, we illustrate how the characteristic features of MtF address specific regulatory requirements and best-practice recommendations.

Mark-to-market trading book rigorously			●
Measure market risk at portfolio level (VaR)	●		●
Apply a comprehensive stress testing program	●	●	
Estimate actual and potential exposures	●	●	
Account for netting and credit enhancement	●	●	●
Aggregate exposures and compare to limits	●	●	●
Adjust VaR for relevant asset liquidity	●	●	
Model funding liquidity (collateral, credit lines, etc.)	●	●	
Measure credit exposures including liquidity risk	●	●	
Stress test the interaction of market, credit and liquidity risk			
Integrate risk model and daily risk management			●
Relate risk model to trading limits			●
Conduct a regular back-testing program		●	
Develop a rigorous stress testing program	●	●	
Calculate VaR at 99th percentile for 10 day horizon	●	●	
Use one year of historical data or more	●		
Include gamma and vega risks in VaR measure	●		
Move towards an explicit simulation over time	●	●	
Option to extend risk model to specific risk	●		
Cover banking and trading book exposures			●
Compute meaningful exposures that include:			
Collateral and other credit enhancements	●	●	
Netting	●		●
Roll-off risk over time	●	●	
Relate exposures to market, credit, liquidity events	●	●	
Aggregate exposures and relate to credit limits			●
Implement proper credit granting processes			●
Create a portfolio credit model for expected and unexpected losses	●	●	●
Move towards risk-based pricing	●		
Introduce risk-adjusted performance measurement	●		
Capture advanced credit products	●	●	●
Stress test the interaction of market, credit and liquidity risk	●	●	
Assess the effect of an economic downturn	●	●	
Model the effect of potential credit events over time	●	●	

2

The Mark-to-Future Methodology

An Introduction

MARK-TO-FUTURE IS AN ENCOMPASSING AND UNIFYING FRAMEWORK FOR ASSESSING FUTURE UNCERTAINTY. IT IS A PARADIGM THAT VALUES FINANCIAL PRODUCTS ALONG MULTIPLE SCENARIOS AT FUTURE TIME STEPS. THE GENERATION OF A MARK-TO-FUTURE VALUE FOR EACH SCENARIO AND TIME STEP REPRESENTS THE COMPUTATIONALLY INTENSIVE FIRST STAGE OF PROCESSING (PRE-CUBE), WHILE THE MEASUREMENT OF PORTFOLIO RISK AND REWARD OCCURS STRICTLY AS A SECONDARY POST-PROCESSING STAGE (POST-CUBE). THIS CHAPTER PROVIDES A GENERAL OVERVIEW OF THE MARK-TO-FUTURE FRAMEWORK.

At the foundation of the MtF framework is the generation of a three-dimensional cube consisting of a series of **MtF tables**. Each MtF table has dimension $S_t \times N$, where S_t is the number of scenarios and N is the number of instruments. While the number of scenarios S_t may vary over time steps (as denoted by the subscript), for ease of exposition the number of scenarios per time step will be assumed constant throughout this document.

Each MtF table is associated with a given time step t over an overall time horizon of T steps. A pre-computed **MtF Cube** provides a basis onto which the mapping of all financial products and all positions in those products can be accommodated, thereby enabling the full characterization of future portfolio distributions for multiple portfolio regimes through time. Figure 1.1 illustrates a representative MtF Cube, which consists of a series of MtF tables, one for each time step t.

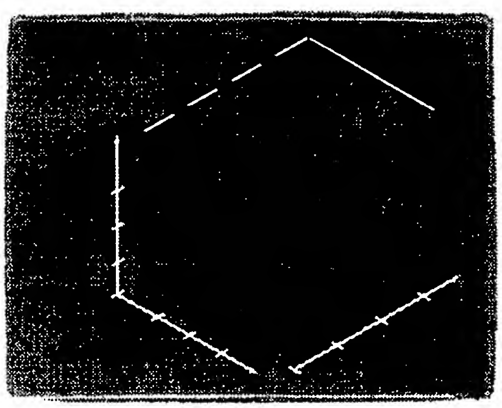

Figure 1.1: Representative MtF Cube

Each cell in the MtF Cube contains the simulated MtF value for a given instrument under a given scenario and time step. In certain applications, other sensitivity measures such as an instrument's delta or its duration may be included in addition to a MtF value. Therefore, in the general case, each cell of a MtF Cube contains a vector of risk-factor-dependent measures for a given instrument under a given scenario and time step. In some applications, the vector may even contain a set of risk-factor dependent cashflows for each scenario and time step. For ease of exposition, however, this document focuses primarily on the typical case where each cell contains only the instrument's MtF value.

Fundamental to the MtF framework is the premise that generating a **Basis MtF Cube** need not depend on the knowledge of portfolio holdings: a single Basis MtF Cube accommodates the risk-reward assessment of multiple **regimes** (portfolio strategies) simultaneously. The representation of any given regime such as a **buy and hold** strategy or an **immunization** strategy is simply captured through the mapping of positions onto the MtF Cube. Thus, many portfolio regimes may be mapped simultaneously onto the same pre-computed MtF Cube.

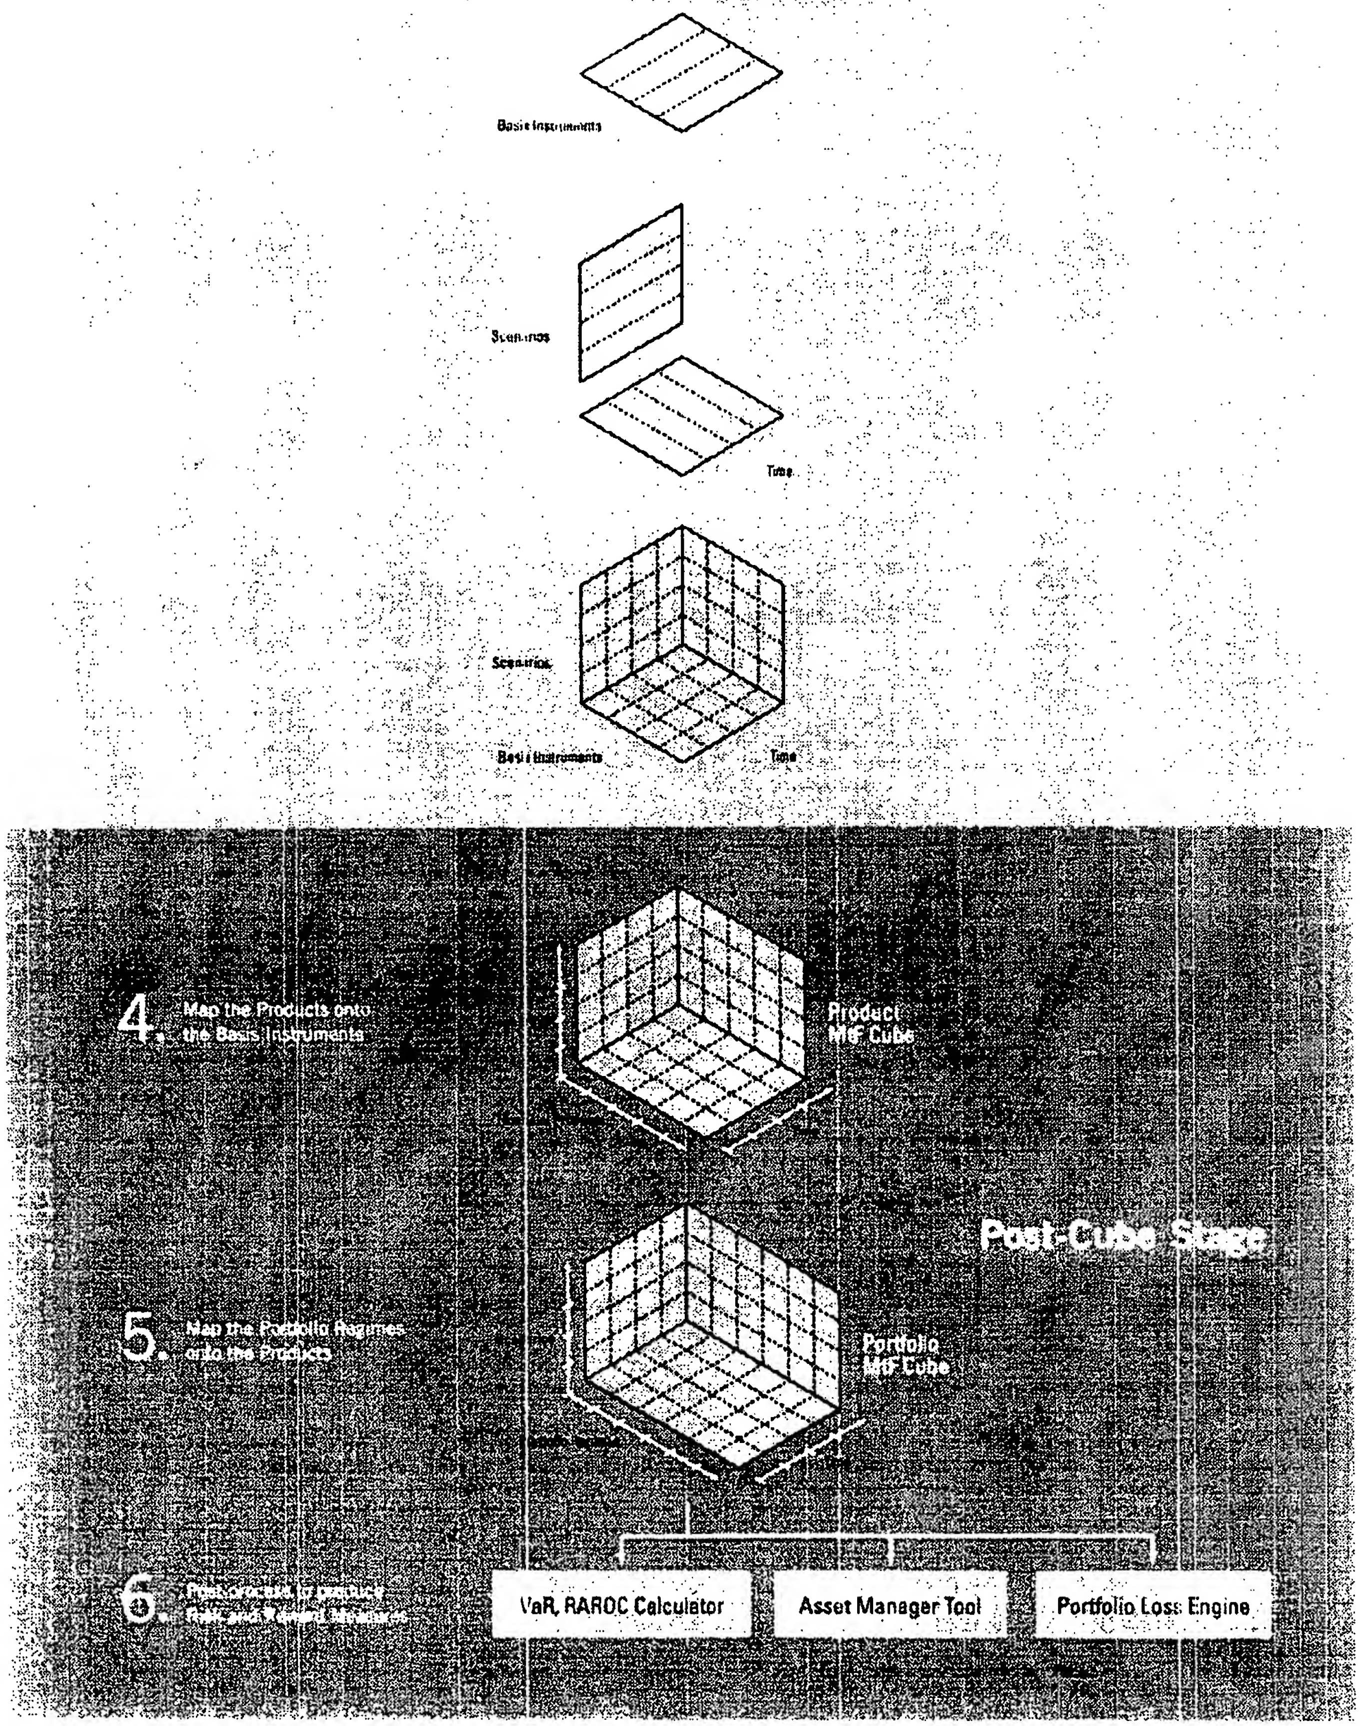

Basis Instruments
Scenarios
Time
Scenarios
Basis Instruments
Time
4. Map the Products onto the Basis Instruments
Product MtF Cube
Post-Cube Stage
5. Map the Portfolio Regimes onto the Products
Portfolio MtF Cube
6. Post-process to produce Risk and Reward Measures
VaR, RAROC Calculator
Asset Manager Tool
Portfolio Loss Engine

Basis: M	$i=1,\ldots,N$ basis instruments $j=1,\ldots,S_t$ scenarios $t=1,\ldots,T$ time steps	m_{ijt}		
Product: M′	$i'=1,\ldots,N'$ financial products $j=1,\ldots,S_t$ scenarios $t=1,\ldots,T$ time steps	$m'_{i'jt}$	$M'_t = f_t(M_t)$	$f_t()$
Portfolio: M″	$i''=1,\ldots,N''$ portfolio regimes $j=1,\ldots,S_t$ scenarios $t=1,\ldots,T$ time steps	$m''_{i''jt}$	$M''_t = M'_t x_t$	$x_{i't}$

Table 1.1: Summary of Mapping and MtF Cube Notation

The MtF approach to risk-reward assessment is summarized by the following six steps, each of which can be explicitly configured as an independent component of the overall process:

1. Define the basis instruments.
2. Define the scenario paths and time steps.
3. Simulate the instruments over scenarios and time steps to generate a Basis MtF Cube.
4. Map the financial products onto the Basis MtF Cube to form a Product MtF Cube.
5. Map the portfolio regimes onto the Product MtF Cube to form a Portfolio MtF Cube.
6. Post-process to produce risk and reward measures.

The generation of the Basis MtF Cube in Steps 1 to 3 of the overall process are the *only* computationally intensive steps and, significantly, need to be performed only *once*. These steps represent the **pre-Cube stage** of MtF processing. In contrast, Steps 4 to 6 represent straightforward mapping exercises, which can be performed in virtual real-time and with minimal additional processing. These three mapping steps represent the **post-Cube stage** of MtF processing. The diagram on the previous page illustrates the process flow incorporating the six key steps of the MtF framework.

The Basis MtF Cube generated in the pre-Cube stage is the only MtF Cube that necessarily exists in a physical sense. The Product and Portfolio MtF Cubes generated in the post-Cube stage are completely defined by the product and portfolio mappings and, thus, typically exist only in a conceptual sense. Table 1.1 provides a summary of the mapping and MtF Cube notation used in this document.

Note that, in certain applications, additional intermediate mapping steps may be incorporated into the post-Cube stage which generate other intermediate (but conceptual) MtF Cubes. An example may include an intermediate step mapping financial products onto *hypothetical* basis instruments that precedes the mapping of the hypothetical instruments into actual basis instruments. Nonetheless, for ease of exposition, this document focuses on the two key mapping steps that produce the Product and Portfolio MtF Cubes, respectively.◎

Define the Basis Instruments

The Basis MtF Cube is generated by simulating a series of basis instruments across predefined scenarios and time steps. **Basis instruments** may represent actual financial products such as stocks and bonds, synthetic instruments such as a series of zero coupon bonds or, in some cases, abstract indices such as the inputs to a multi-factor model.

The choice of basis instruments to be simulated in the pre-Cube stage of the MtF framework depends on the overall business application and may involve trade-offs between the magnitude of data storage (the dimensions of the Basis MtF Cube) and pricing accuracy. Nonetheless, a significant amount of dimensionality reduction can often be achieved with little or no loss of pricing accuracy by the judicious choice of basis instruments.

In most applications, the Basis MtF Cube will consist of a mixture of actual products representing exchange-traded securities or OTC securities held in inventory and synthetic instruments onto which newly transacted OTC securities may be mapped. An understanding of how the values of risk factors determine product valuation helps in the selection of basis instruments. The MtF value, $m'_{i'jt}$, of financial product, i' ($i' = 1,\ldots,N'$), under scenario, j ($j = 1,\ldots, S_t$), at time step, t ($t = 1,\ldots,T$), is a function of a number of risk factors, $u_k = u_1,\ldots, u_K$, and, in general, can be represented as

$$m'_{i'jt} = g'(u_1, u_2, u_3,\ldots, u_K)$$

The MtF value, m_{ijt}, of basis instrument i ($i = 1,\ldots,N$) under scenario j at time step t is also a function $g()$ of a number of risk factors and similarly can be represented as

$$m_{ijt} = g(u_1, u_2, u_3,\ldots, u_K)$$

When the basis instrument represents a synthetic instrument or an abstract index, it is often a function of a single risk factor. Thus, in the typical case

$$m_{ijt} = g(u_k)$$

for risk factor k.

The key benefit derived from the use of basis instruments is the ability to store all or a portion of the computation of the product MtF values in the Basis MtF Cube. Therefore, the MtF value of a given financial product becomes a function, $f()$, of the pre-computed MtF values of the basis instruments or simply

$$m'_{i'jt} = f(m_{1jt}, m_{2jt}, m_{3jt},\ldots, m_{Njt})$$

In the case where basis instruments are actual financial products, $m_{i'jt}' = m_{ijt}$, and the mapping of the financial product onto the basis instrument is a straightforward

one-to-one mapping. In contrast, when the basis instruments are synthetic instruments, the mapping of the financial product becomes more complex.

Determining the ideal composition of basis instruments, therefore, requires the *a priori* definition of a set of functions mapping financial products to basis instruments and ultimately to risk factors. A **mapping sequence**, as illustrated in Figure 2.1, represents a configuration of risk factors and basis instrument mappings used to value a given financial product.

Figure 2.1: Mapping sequence of products onto instruments onto risk factors

It is necessary to define a mapping sequence in order to select the appropriate basis instruments to be contained in the Basis MtF Cube. The next sections describe some of the issues involved in defining a mapping sequence. Issues surrounding the mapping functions are discussed in Step 3.

Mapping Sequence Considerations

A variety of possible mapping sequences may be defined. These possibilities are based on the modeling choices that are made prior to generating the Basis MtF Cube. As an example, consider the mapping sequence A chosen to generate a product MtF cube containing a non-dividend paying stock and a forward contract, with expiry t, on that stock. Mapping sequence A is illustrated in Figure 2.2.

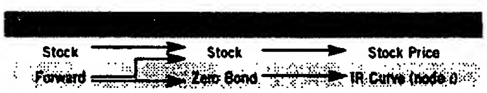

Figure 2.2: Mapping sequence A

In this particular sequence, the mapping of the stock onto the basis instrument is one-to-one. The MtF value of the basis instrument is modeled as a simple realization of the risk factor itself under each simulation. The forward contract is mapped onto two basis instruments: the stock and the zero coupon bond.

Consider a second mapping sequence B with the stock, in this case, simulated by a generic two-factor equity model where the MtF values of the stock are determined as a function of the stock's specified factor loadings onto the equity factors. This mapping sequence is illustrated in Figure 2.3.

Figure 2.3: Mapping sequence B

In this sequence, the mapping of the stock onto the basis instrument is still one-to-one, but the MtF value of the basis instrument is based upon the generic two-factor equity model.

Finally, consider a third mapping sequence C that generates a Basis MtF Cube using the equity factors as the basis instruments. In this case, the MtF values of the stock are determined as a function of the factor loadings in the equity factor basis instruments. This mapping sequence is illustrated in Figure 2.4.

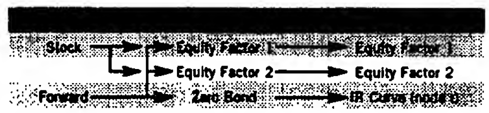

Figure 2.4: Mapping sequence C

In this case, the stock is mapped onto two equity factor basis instruments. The number of basis instruments required using mapping sequence C is always two as all stocks map onto the same two risk factors, regardless of the number of stocks in a portfolio. Only the mappings (i.e. the factor loadings) associated with each equity are product-specific.

While it is interesting to compare alternative modeling techniques, the choice of mapping sequence often involves other, more pragmatic, trade-offs. For example, mapping sequence A may be preferred when full explanatory power is desired for future equity distributions (perhaps based upon historical realizations). However, this approach may prove deficient when newly issued securities (that possess no price history) are addressed.

In contrast, mapping sequence C may be preferred for purposes of minimizing the dimensions of the Basis MtF Cube as well as for mapping equities with no previous price history. However, this approach may be less appropriate when it is important to consider the specific risks associated with individual equities. Note that specific risks can, in fact, be addressed in this particular mapping sequence through a more complex product mapping that directly incorporates the specific risk of the equity as an additional input.

Mapping sequence B may emerge as a preferred alternative in this case. Specific risk and newly issued equities may be addressed easily using this approach since each security is simulated uniquely using the factor model, thereby allowing for the straightforward incorporation of a random component as the stock is

Stock Instrument	s_0	s_1	s_2	s_3	s_4
Settlement Account	0	c_1	$c_1 f_{1,2}$	$c_3 + c_1 f_{1,2} f_{2,3}$	$c_3 f_{3,4} + c_1 f_{1,2} f_{2,3} f_{3,4}$
Bundled Instrument	s_0	$s_1 + c_1$	$s_2 + c_1 f_{1,2}$	$s_3 + c_3 + c_1 f_{1,2} f_{2,3}$	$s_4 + c_3 f_{3,4} + c_1 f_{1,2} f_{2,3} f_{3,4}$

Table 2.1: MtF values of bundled stock instrument

(s_t and c_t represent the value of the basis equity instrument and any cash dividends paid out at time step t along a given scenario. $f_{t-x,t}$ represents the future value factor from time step $t-x$ to t, along a given scenario. The MtF value of the bundled instrument is the sum of the MtF values of the standard instrument and the cash account.)

Marked-to-Future. This has implications for the size of the Basis MtF Cube, which must incorporate MtF values for all individual equities.

Settlement and Reinvestment Considerations

Settlement and reinvestment is another key issue to consider when selecting the appropriate basis instruments. The MtF framework is inherently forward-looking in nature. MtF values are explicitly calculated at future time steps. This requires the appropriate modeling of financial product aging.

Significantly, any settlement and subsequent reinvestment of financial products prior to an overall time horizon must be directly incorporated in the pre-computed Basis MtF Cube. This can only be achieved by defining basis instruments which effectively represent **bundled instrument** positions consisting of a standard basis instrument along with a settlement account that reinvests any cash settled by the basis instrument through time. The bundled basis instrument captures the total return of a financial product through time.

As an example, consider a bundled stock instrument which consists of the basis stock instrument along with a settlement account that captures and reinvests dividends paid out at time steps $t=1$ and $t=3$. Table 2.1 illustrates the MtF values of the bundled basis instrument across a given scenario path for a $T=4$ time horizon.

Figure 2.5 provides an illustration of MtF values for the bundled equity instrument as simulated across $S_j=100$ scenario paths over the $T=4$ time horizon. At each time step, the MtF values capture *both* the distribution of stock prices and the distribution of reinvested dividends.

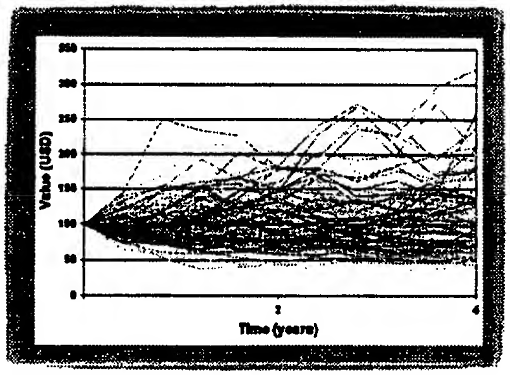

Figure 2.5: MtF values of bundled stock instrument across scenarios

The vertical axis in Figure 2.5 represents MtF value while the horizontal axis represents time. Each node on a line corresponds to the realization of the bundled instrument's MtF value under a given scenario and time step.

As a second example, consider a **bundled zero coupon instrument** consisting of a standard zero coupon bond (with a notional of one and maturity at $t=2$) along with a settlement account which captures and reinvests the notional amount paid out at the maturity date. Table 2.2 summarizes the MtF values of the bundled basis instrument across a given scenario path for the $T=4$ time horizon.

Zero Coupon Bond	$d_{0,2}$	$d_{1,2}$	1	0	0
Settlement Account	0	0	0	$f_{2,3}$	$f_{2,3} f_{3,4}$
Bundled Instrument	$d_{0,2}$	$d_{1,2}$	1	$f_{2,3}$	$f_{2,3} f_{3,4}$

Table 2.2: MtF values of a bundled $t = 2$ zero coupon instrument

($d_{t,t+x}$ and $f_{t-x,t}$ represent the discount factor from time step t to $t+x$ and the future value factor from time step $t-x$ to t, respectively, along a given scenario. The MtF value of the bundled instrument is the sum of the MtF values of the standard instrument and the cash account.)

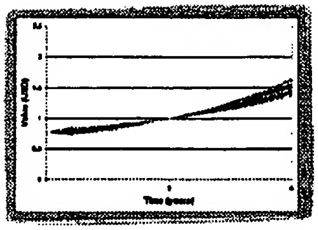

Figure 2.6: MtF values of a bundled $t=2$ zero coupon instrument across scenarios

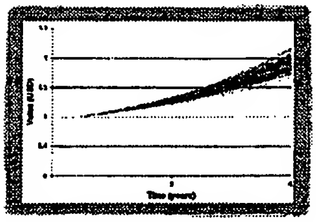

Figure 2.7: MtF values of a bundled $t=0$ zero coupon instrument across scenarios

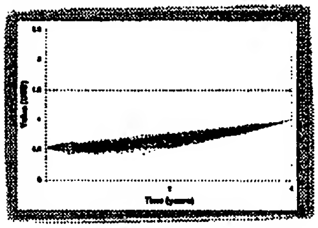

Figure 2.8: MtF values of a bundled $t=4$ zero coupon instrument across scenarios

Figure 2.6 provides an illustration of MtF values for the bundled zero coupon instrument with maturity of $t=2$ as simulated across $S_4=100$ scenario paths over the $T=4$ time horizon.

In comparison, Figures 2.7 and 2.8 are illustrations of two additional bundled zero coupon basis instruments (for maturities of $t=0$ and $t=4$, respectively) with MtF values simulated across $S_4=100$ scenarios over the $T=4$ time horizon.

Note that the $t=4$ bundled instrument poses only price risk over the time period $t<4$, during which its MtF value is sensitive to interest rate changes. At its $t=4$ maturity date, when the notional is settled, it represents a riskless position. In contrast, the $t=0$ bundled instrument is riskless at its $t=0$ maturity date, but poses reinvestment risk over the duration of the $T=4$ time horizon as the notional amount is reinvested. The $t=2$ bundled instrument has a changing risk profile over the $T=4$ time horizon; it poses price risk over the period $t<2$, is riskless at $t=2$ and poses reinvestment risk over the period $2<t\leq 4$. These examples illustrate the significance of considering settlement and reinvestment risk through time even for a straightforward instrument such as a zero coupon bond.

In a more complex example, a bundled instrument may consist of a standard basis instrument, one or more underlying instruments, as well as a settlement account. Tables 2.3 and 2.4 describe a bundled call option on a bond future consisting of an exchange-traded bond futures option instrument (expiry at $t=1$), an exchange-traded bond future instrument (expiry at $t=3$), an exchange-traded bond instrument, and in addition, a settlement account.

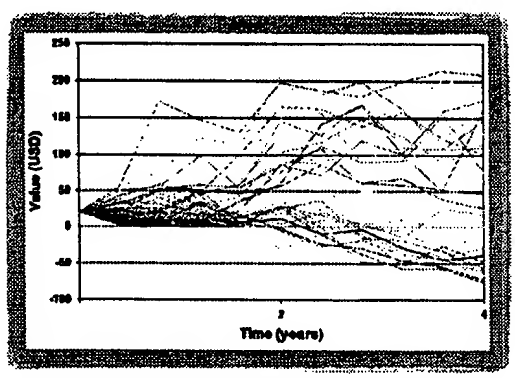

Figure 2.8: MtF values for a bundled bond futures option across scenarios

Bond Futures Option	C_0	$C_1=F_1-K_c>0$ (in-the-money)	0	0	0
Bond Future Instrument	0	0	F_2-F_1	F_3-F_2	0
Bond Instrument	0	0	0	0	B_4
Settlement Account	0	0	C_1f_{12}	$C_1f_{12}f_{2,3}+(F_2-F_1)f_{2,3}$	$C_1f_{12}f_{23}f_{34}+$ $+(F_3-F_2)f_{3,4}-F_3f_{3,4}$
Bundled Instrument	C_0	C_1	$F_2-F_1+C_1f_{12}$	F_3-F_2+ $(F_2-F_1)f_{2,3}+C_1f_{12}f_{23}$	$B_4-F_3f_{3,4}+(F_2-F_1)f_{2,3}f_{3,4}$ $+(F_3-F_2)f_{3,4}+C_1f_{12}f_{23}f_{34}$

Table 2.3: In-the-money MtF values of a bundled bond futures option

(C_t, F_t and B_t represent the bond futures call option, bond future and bond (all-in) prices, respectively, at time step t along a given scenario. K_c represents the option strike price. $f_{t-x,t}$ represents the future value factor from time step $t-x$ to t, along a given scenario where the call is in-the-money at $t=1$. The MtF value of the bundled instrument is the sum of the MtF values of the standard instrument and the cash account.)

Bond Futures Option	C_0	0 (out-of-the-money)	0	0	0
Bond Future Instrument	0	0	0	0	0
Bond Instrument	0	0	0	0	0
Settlement Account	0	0	0	0	0
Bundled Instrument	C_0	0	0	0	0

Table 2.4: Out-of-the-money MtF values for a bundled bond futures option

(C_t, F_t and B_t represent the bond futures call option, bond future and bond (all-in) prices, respectively, at time step t along a given scenario. K_c represents the option strike price. $f_{t-x,t}$ represents the future value factor from time step $t-x$ to t, along a given scenario where the call is out-of-the-money at $t=1$. The MtF value of the bundled instrument is the sum of the MtF values of the standard instrument and the cash account.)

Figure 2.9 provides an illustration of the MtF values of a bundled bond futures call option instrument (with option expiry of $t=1$ and futures expiry of $t=2$ as simulated across $S_4=100$ scenarios over the $T=4$ time horizon). Along some scenario paths the call option expires in-the-money at $t=1$ and, hence, settles onto the bond future, with the future itself settling onto the cheapest-to-deliver bond at $t=2$. Along other scenario paths the call option expires out-of-the-money at $t=1$ and, in these cases, the entire bundled position is worthless from that point on.

This example demonstrates the importance of appropriately incorporating settlement and reinvestment over time when physical settlement is a key component of product aging. ⊚

Define Scenario Paths and Time Steps

Future distributions of portfolio outcomes are ultimately driven by future distributions of underlying risk factors such as interest rates, equity indices or exchange rates. In traditional risk-reward methodologies, the selection of scenarios is not arbitrary; rather, a very specific risk factor distribution is implied. A summary statistic such as a covariance matrix of factor returns that describes the risk factor distribution typically serves as the primary input into the analysis. However, only for very specific distributions can a small number of parameters for each factor truly describe the distribution. As a consequence, the use of a specific summary statistic as the key input (with no additional knowledge of the true underlying distribution) severely constrains an analysis when the underlying distribution is not well described by that statistic.

In contrast, in the MtF framework, the scenario choice is a key input parameter to the overall risk assessment process. Since scenarios are the primary inputs, the distribution of underlying risk factors is defined by the explicit set of scenarios actually chosen. As such, any desired future risk factor distribution can be directly incorporated into the analysis in a manner that is not constrained by a specific summary statistic.

Scenarios are the language of risk in a MtF framework. In practice, the appropriate choice of scenarios, whether explicit or implicit, is the key factor that determines whether or not a risk analysis is adequate. Likewise, the quality of a MtF simulation depends on the ability to generate relevant, forward-looking scenarios that properly represent the future. This chapter discusses the value of scenarios, outlines the benefits of using scenarios and considers several sources of scenarios including historical data, scenario proxies and scenario bootstrapping.

What are Scenarios?

A **pricing model** determines the value of a financial instrument as a function of the values of a number of instrument attributes and risk factors. A complete set of values for these risk factors is called a scenario. A **scenario** describes the evolution of the values of the risk factors over time. The specification of a scenario at a given future point in time describes a possible state-of-the-world at that time. The current state-of-the-world, described by the current market data, is called the **nominal scenario**. The nominal scenario is used to produce the mark-to-market value; MtF scenarios are used to calculate the value of an instrument at a future time. There are several different types of scenarios related to market, liquidity and credit risk as well as joint scenarios that cover a combination of these risks.

Market risk scenarios are described by market risk factors. Standard market risk factors include various

interest rate term structures, equity indices, exchange rates, commodity prices or implied volatilities. Credit risk scenarios are described by specifying credit risk factors. For example, Credit Metrics (Longerstaey and Zangari 1996) defines country, region and sector indices as credit drivers for the asset value and credit quality of each obligor. Alternatively, CreditRisk+ (Credit Suisse 1997) specifies default probabilities directly as the key credit risk factors. Liquidity risk scenarios are described by risk factors such as bid-ask spreads and market depth or trading volumes.

Joint scenarios on market and credit risk might be described by specifying market and credit risk factors in the same scenario. A model that integrates market and credit risk factors is presented by Iscoe et al. (1999). Consider a corporate bond. The present value of the bond is a function of the Treasury curve and an additional credit spread that compensates bond holders for the default risk of the corporate issuer. The credit spread, in turn, depends on the credit quality of the corporate issuer. If a credit downgrade occurs, cash flows must be valued using a higher credit spread. Scenarios that include interrelated outcomes for the issuer's credit rating and for the Treasury curve link market and credit risk consistently.

Joint scenarios on market and liquidity risk for equity portfolios might include equity indices and specific risk factors to capture market risk and equity-specific, simulated holding periods to capture liquidity risk. Joint scenarios on market, credit and liquidity risk for emerging market bonds, for example, may cover US Treasury rates, credit risk factors such as credit drivers, sovereign spreads or default probabilities and bid-ask spreads.

The Value of Scenarios

A **MtF value** is the value of a financial instrument at a future time under a given scenario. There is no need to consider the likelihood or probability of that scenario when computing the MtF value. Some risk-reward measures (e.g., Regret (Dembo 1991, 1999; Dembo and Freeman 1998) and Average Shortfall (Artzner et al. 1998)) are based on probability-weighted MtF values. Value-at-Risk is the MtF value at a specific point in the tail of a MtF distribution. At a later stage (Step 6) when risk/reward statistics are computed, it is necessary to assign probabilities to these scenarios, but an important feature of the MtF approach is that the choice of scenarios need not be dependent on the probability assigned to them. This feature allows MtF values to be computed *once* and subsequently used as inputs into many different risk measures.

Regulators are calling for an integrated measure of market, credit and liquidity risk. In a MtF framework, these disparate sources of risk are linked quite naturally at the scenario level, without reference to the risk measure. Scenarios that embody correlated, consistent, simultaneous changes in market, credit and liquidity states naturally provide correlated, consistent MtF output, from which risk measures that link these sources of risk can be calculated. Various integrated measures of risk and return can be obtained from this consistent MtF output. Thus, MtF is a unifying framework that simplifies integration of the measurement of risk and reward for capital allocation purposes.

The use of scenarios to describe the future possible states-of-the-world makes the risk management process transparent, descriptive and manageable. Non-technical managers can understand a set of scenarios and the market intuition of these managers can be exploited to define them.

Requirements of Scenarios

Scenarios provide significant benefits in risk management, but in order to do so, they must have certain characteristics. The scenario choice must be explicit, the scenarios chosen must span all relevant future outcomes and the relevant time horizon or horizons, and the scenarios must be forward-looking.

Ultimately, the MtF value of a portfolio is a function of the values of the risk factors constituting the scenarios. However, small market moves may cause large changes in the value of a portfolio of instruments with non-linear pay-offs. Thus, the distribution of outcomes in the risk factor or scenario space may be very different from the distribution of outcomes in the portfolio space. The impact of an individual scenario cannot be determined until after the MtF value is computed. It is a misconception, for example, to think that stressful market events such as the crash of '87 are necessarily stressful scenarios.

The scenarios in a MtF simulation must span a wide range of possible outcomes, without regard to the likelihood of those outcomes. A distinction is often made between extreme scenarios that are used for stress testing and scenarios based on typical market conditions. The MtF framework does not distinguish between them. All scenarios are part of the same consistent MtF Cube. A distinction is made only at the post-Cube stage of MtF processing.

For the purposes of risk management, knowing the outcomes under extremes is essential. A fundamental function of risk management is to decide whether or not to hedge these extreme events and at what expense. Measures such as regret and average shortfall capture the

shape of the entire downside of the distribution; omitting these extremes affects the shape of the downside and thus the value of the measure. In the MtF framework, measures such as put value and call value place a value or cost on insuring the downside and upside of a given portfolio. The put value and call value measures are defined in Step 6.

Scenario selection differs considerably from forecasting. A forecast is a prediction that a single scenario will occur. The accuracy of a forecast is therefore crucial, yet as noted by John Lipsky, the Chief Economist at Chase Manhattan, "No one is able to consistently foresee specific market developments far into the future." (Fuerbringer 1999). The goal in selecting scenarios in a MtF analysis is to span the range of future events, not to forecast that any of these events will actually occur.

The scenarios must span a wide range of possible outcomes and they must also extend over a horizon, or multiple horizons, of appropriate length. This requires the generation of scenarios over time (Figure 3.1). The appropriate horizon for a market risk measure might be one day, 10 days or longer, depending on the liquidity of the position. Estimating exposure profiles for counterparty risk may require multiple time horizons over 10 years or more.

Liquidity risk measurement depends on being able to measure the risk of portfolios that evolve over time. Though the cost of liquidity is usually reflected in wide bid-offer spreads, this is difficult to model directly. Liquidity risk can be modeled indirectly, by designing joint scenarios on daily trading volumes and market risk factors such as price and volatility. In this model (Yung 1999c), portfolio positions are liquidated conditional on the simulated outcomes for these risk factors in each scenario. For example, when simulated trading volumes are high, positions will be liquidated faster than when trading volumes are low.

In order to interpret a risk measure, consumers of risk information, such as senior managers and boards of directors of financial institutions, must understand the nature of the scenarios used to produce the reports they read. Scenarios are the critical elements that define whether or not a risk analysis is adequate. MtF requires an explicit scenario choice. The end users of the risk information should help sanction the scenarios used because scenarios are their vehicle for understanding and interpreting the output of the simulation. Explicit knowledge of scenarios is essential for good risk-reward management. The good news is that anyone experienced in the market, who is involved in taking or managing risk, regardless of their technical knowledge, is capable of understanding and generating interesting and relevant scenarios.

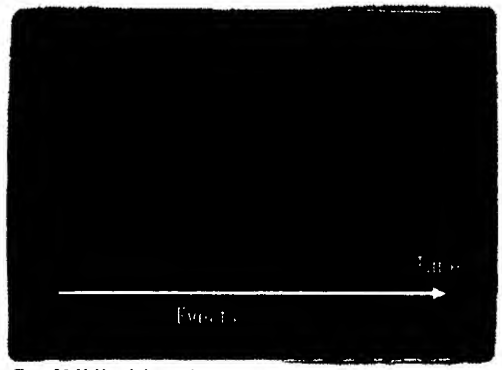

Figure 3.1: Multi-period scenarios

The scenarios used are key in determining whether a MtF calculation is forward-looking. To be truly forward-looking, scenarios must contain events that are not in the current information set. This seems somewhat paradoxical. How do we define scenarios that are not part of the current information set? Forward-looking views can be extracted from proxy events or by other techniques, such as scenario bootstrapping. Scenario proxies and bootstrapping are discussed following a review of traditional approaches to risk measurement.

What is the Value-at-Risk of a position of 300,000 shares of Gold Fields, the second largest gold producer in South Africa? Standard Value-at-Risk analysis at a 95% confidence interval indicates a probable loss of 5%, assuming that the entire position can be liquidated in the next 24 hours. Given that the average daily trading volume of Gold Fields on the Johannesburg Stock Exchange is only 175,000 shares, portfolio VaR is increased due to reduced liquidity–but by how much?

A Mark-to-Future Model

Scenarios on Gold Fields' market price and daily trading volume describe the dynamic nature of liquidity risk:

1. **Market Risk Scenario Set:** 100 one-day shock scenarios on share price. This scenario set does not account for the impact of liquidity.
2. **High Volatility Scenario Set:** In general, as market volatility increases, liquidity dries up. Five scenarios at multiple time horizons are generated based on the trading volumes in the five most volatile periods in Gold Fields' share price history.
3. **Low Trading Volume Scenario Set:** The impact of extremely low trading volumes is captured by bootstrapping 10 scenarios at multiple time horizons based on the 250 lowest trading volumes from Gold Fields' recent history.

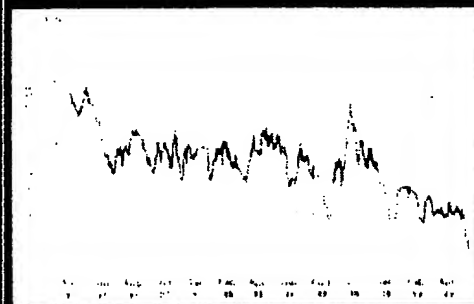

The fall of Gold Fields (Share price = 100 on April 14, 1997)

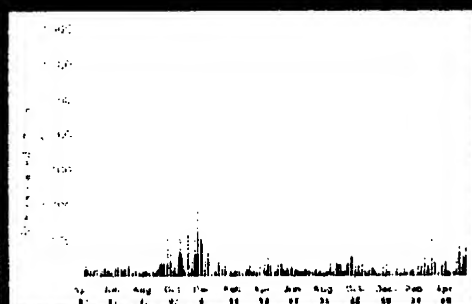

Gold Fields' Daily volume

The cost of liquidity and the time to liquidate is modeled by liquidation strategies. A strategy describes how the position is unwound; proceeds from the disposition are settled into cash.

1. **Instantaneous Liquidation:** This strategy assumes that infinite liquidity exists in the market. All 300,000 shares of Gold Fields are liquidated at the quoted market price in one day.
2. **Unconditional Liquidation:** 10% of the Gold Fields shares are sold each day over the next 10 days at the quoted market price.
3. **Conditional Liquidation:** Sale of Gold Fields is restricted to 20% of each day's trading volume, assuming that at this rate, shares can be sold at the quoted market price.

Scenarios and liquidation strategies are specified for each of the four liquidity categories: infinite, average, limited and negligible liquidity.

	Simulation Category	Scenario	Liquidation Strategy
Infinite Liquidity	[illegible]	[illegible]	[illegible]
Average Liquidity	[illegible]	Market Risk Scenario [illegible]	[illegible]
Limited Liquidity	[illegible]	[illegible]	[illegible]
Negligible Liquidity	[illegible]	Market Risk Scenario and Low Trading Volume Scenario	Conditional Liquidation

Infinite Liquidity	5.2	5.2	-	1 day
Average Liquidity	12.3	5.2	7.1	10 days
Limited Liquidity	13.2	5.2	8.0	18 days
Negligible Liquidity	21.0	5.2	15.8	36 days

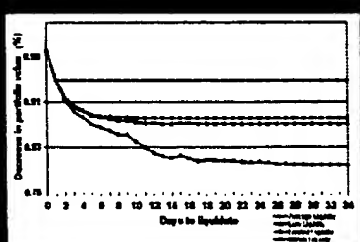

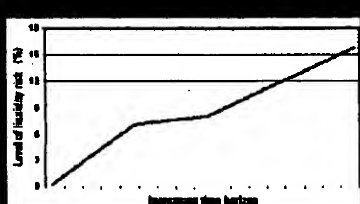

Traditional Approaches

In some methodologies, namely covariance-based analytical methods, the scenarios are implicit in the choice of historical data used to build the covariance matrix and are therefore obscured. Thus, the specification of the scenarios is hardcoded into the calculation of the risk measure. It is rare that the consumers of risk information (senior managers in financial institutions, the board, etc.) know the nature of the scenarios that are implied in the reported risk measures. Except in extreme cases, very few people can determine if a covariance matrix is reasonable.

Structured Monte Carlo methods generate risk factor scenarios from a covariance matrix. The scenarios that are implied in the covariance matrix are recovered and stated explicitly. However, the perceived richness of a scenario set generated in a Monte Carlo simulation model often masks the fixed modeling choices that went into estimating the matrix.

For example, methods currently in use to estimate covariance for risk measurement may weight past events differently. The Bank for International Settlements (BIS) accepts scenarios that are obtained by looking at one year of historical, daily outcomes equally weighted (BIS 1996). The RiskMetrics approach is based on an exponentially weighted history with a 6% daily decay factor, which translates into a small probability for scenarios that are more than 60 days old (Longerstaey and Zangari 1996).

Each of these methods corresponds to a different choice of scenarios and differing assessments for the way in which probabilities are assigned to the scenarios. For example, using the scenarios suggested by RiskMetrics implicitly assumes that the last 60 days includes all the fundamental shifts in the market that are likely to occur between now and the risk horizon. Furthermore, their weighting suggests that the only important events have occurred in very recent history; events of the past are considered very unlikely.

Normal distributions often understate the probability of extreme events; historically observed distributions have fat tails when compared to most parametric distributions.

Consider a particular example. In the last 84 years, there have been nine crashes in the Dow Jones Industrial Index whose daily falls exceeded six standard deviations (based upon the previous 250 business days). The likelihood of nine crashes of that magnitude occurring in 84 years is 0.04%. Thus, a scenario of this magnitude is effectively ignored under a traditional approach.

Neglecting the possibility of a crash scenario is clearly not prudent in a comprehensive risk analysis. Yet, if only recent history is considered, or the model is restricted to normal distributions that do not include fat tails, this

possibility may be neglected. Under the MtF framework, a crash scenario can be explicitly incorporated. Risk analyses can include both historical scenarios to account for typical events and stress test scenarios to account for atypical events. In this case, one might consider all of the developed world's stock market history and include the large market movements that were observed as possible scenarios for the future.

Consider a second example. At the end of March 1998, the Canadian dollar was trading at approximately 0.708 USD. The Canadian dollar had last touched 0.700 USD in 1994. Traditional risk measures, based on a summary statistic such as standard deviation, applied to a portfolio with Canadian dollar exposure would have provided a false sense of security to the holders of that portfolio. Even if the standard deviation had been based on one year of CAD/USD history, the implied set of scenarios underlying this statistic could have never gone beyond the bounds of that history. In fact, using this approach over a six month horizon, the drop in the Canadian dollar from 0.708 to 0.633 USD that occurred between March and August 1998 represented a three standard deviation move with a 0.02% likelihood of occurring. (The standard deviation of daily log returns with zero mean is calculated based upon one year of daily observations between March 1997 and March 1998 and scaled to a five month horizon by the square root of time).

Sources of Forward-Looking Scenarios

There are many ways to generate scenarios. With some careful analysis, and a broad view on where to seek information for scenario generation, it is possible to find reasonable, explicit, forward-looking scenarios that span the range of possible future states-of-the-world and the time horizon relevant for risk-reward measurement.

The definition of comprehensive and forward-looking scenarios is rooted in history. However, history can be applied creatively. First, the history of each risk factor can be augmented by including more distant events from the risk factor's own history than would be a part of a standard fixed calibration period. Second, the history of one risk factor can become a proxy for scenarios on another factor.

Using History

History enables us to generate reasonable sequences of events. Traditional risk-reward approaches often use a historical time series as the sole indicator of future risk factor distributions. Historical time series over long periods capture a very wide range of scenarios. Since these events have actually occurred in the market, it is hard to reject them as possible candidates for future events. However, since significant events occur infrequently, a long history may be required to produce a set of scenarios that span the possible range of outcomes. The longer the time series, the wider the range of possible outcomes that can be taken into account.

In established markets where no fundamental shifts in the underlying structure have recently occurred and there is a very long history, history may be an adequate source of scenarios describing typical market conditions. However, the future is sure to contain events outside this information set. When history is the sole source for scenarios, it is implicitly assumed that between now and the horizon there will be no fundamental shift in the underlying forces on the market, other than those already observed during the historical period chosen. This precludes the incorporation of novel or forward-looking views that could have been extracted from other events or by other techniques.

Beyond that, how would we understand the risk or return of a stock that has not yet been issued or of the Euro before it became a traded currency? How can we measure the risk of Internet stocks when there is insufficient history in this market on which to base historical scenarios?

In situations such as these we must extend history and generate future events that reasonably span the range of possible outcomes. One way is to use the history for similar situations in different markets as a proxy.

In January 1999, US-based investors were keenly interested in understanding how events in Brazil might unfold should the real be allowed to float: devaluation of the real by the central bank would cause the rates on Brazilian Brady bonds to soar. Just as the ruble devaluation was the catalyst for market turbulence in August 1998, might pressure building up behind the real have a similar impact in January 1999?

Devaluation of the ruble (August 14-28 1998)

US 30-year Treasury rate (August 14-28, 1998)

Brazilian Brady spreads (August 14-28, 1998)

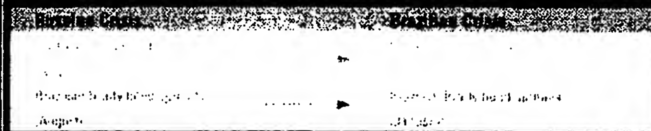

Parallels to Russia

In August 1998, Russia announced the devaluation of its currency and temporary default on its government debt. Russian stocks fell by more than 35% while the ruble tumbled by more than 50%.

Yields on emerging market debt soared, including those on Brazilian Brady bonds, while those on US Treasury bonds reached historic lows. Spreads on the Russian Ministry of Finance Veneshecononbank bonds doubled from their usual norm and the US Treasury zero curve dropped by more than 6% from August 14 to 28. The flight to quality that commenced as investors rebalanced their portfolios caused liquidity in emerging markets to evaporate.

Once confidence in Russia waned, capital flight became precipitous and financial markets reacted fervently. Credit spreads widened, equity markets declined, volatility increased and bid/offer spreads on emerging market debt drifted apart.

History Repeats Itself

Using the period of the ruble devaluation as a proxy for current market scenarios, the two weeks from August 14 to 28, 1998 become two weeks in January 1999. the US Treasury zero rate and the Brazilian Brady bond spread for those 10 business days are mapped to the US rate and the Brazilian Brady spread (respectively) from January 11 to January 25.

Compared to the 95% VaR estimate, the extreme event proxied by the Russian scenario confirms suspicions of the magnitude of the potential risk– it forecasts a 57% decline in portfolio value over this two-week period.

This estimate foreshadows the actual behaviour of the Brady portfolio. Brazil experienced a de facto devaluation of its currency on January 12 when the central bank stopped defending the real. On January 15, 1999, the Brazilian government let the real float and raised interest rates; investors' relief with the government's responsiveness to the crisis is evident from the actual behaviour of the Brady portfolio. This sentiment was comparable to that felt by investors when Russia announced the details of its debt-restructuring program on August 24. ©

Evolution of portfolio value (January 11-26, 1999)

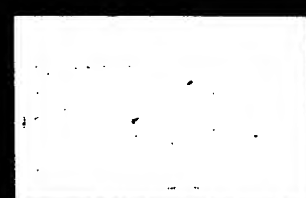

Impact of government intervention (January 15-26, 1999)

The impressive 34% return of the S&P 500 since January 1998 pales in comparison to the performance of Internet stocks. The AMEX Internet Index, which tracks 50 such stocks trading on the New York Stock Exchange, rose by more than 200% during the same period. Gains experienced by blue-chip Internet stocks were even more significant. Overall, share values of companies like Amazon.com, Yahoo!, eBay and AOL have grown by more than 1,000% since the start of 1998. Internet stocks have experienced exceptional growth. A position of five million shares in Amazon.com purchased in early 1988 was worth only $65 million; a year later that position was worth $700 million.

Have previous markets enjoyed growth similar to Internet stocks today? When Nokia first entered the telecommunications industry in the US, its stock enjoyed growth akin to what Internet stocks are experiencing today: the Nokia A share trading in the NYSE has risen by more than 20 times its value since January 1994. The gold rush in the late 1970s also has uncanny similarities, when the US Federal Reserve Board reduced interest rates in response to a threatened debt default by Mexico in 1979, gold took off, rising 300% in six months.

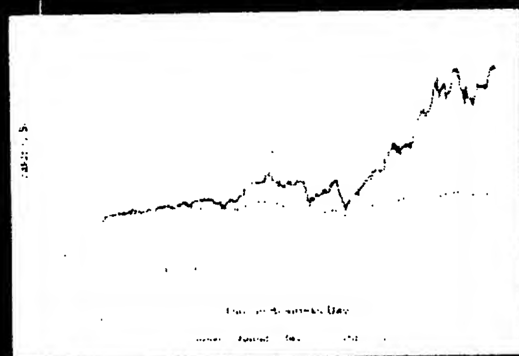

AMEX Internet Index vs. S&P 500 (Jan. 1998 - Mar. 1999)

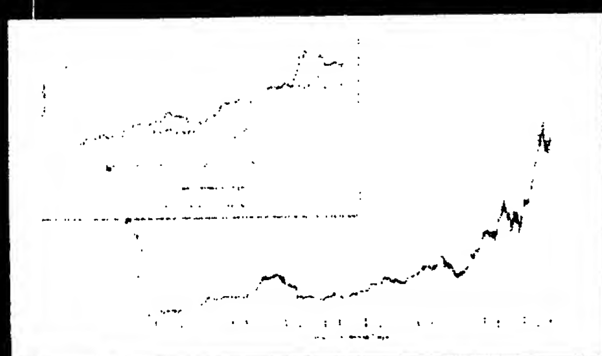

Nokia A shares (Jan. 1994 - Mar. 1999) vs. AMEX Internet Index (Mar. 1998 - Mar. 1999)
Inset: Comparison of Nokia A in 1994 and AMEX Internet Index in 1998

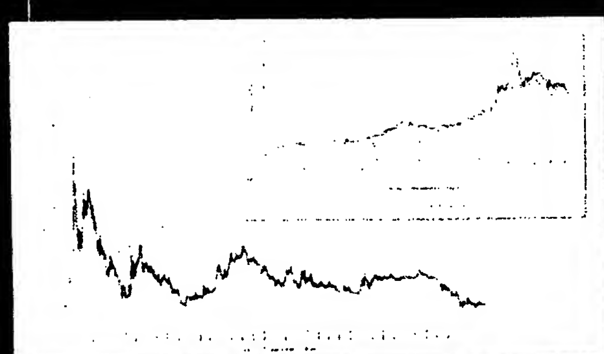

Gold prices (June 1979 - Mar. 1999) vs. AMEX Internet Index (May 1998 - Mar. 1999)
Inset: Comparison of Gold (1979-1980) to Amex Internet Index (1998 - 1999)

However, these two historical episodes had dramatically different outcomes. While Nokia A continues to reach new heights, gold reached its peak in January 1980. To date, gold has fallen by more than 65% from its high with little hope of regaining its original lustre. The crash in the gold market also compares with the impact of the 1929 Depression in the US, Black Monday of October 1987 and the decline in the Japanese stock market in 1989.

1929 Stock Market Crash
(Dow Jones Industrial 1925–1935)

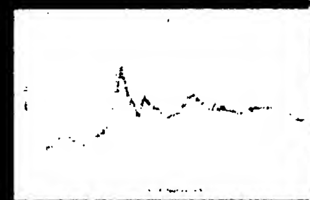

1980 Gold Crash
(Gold prices 1973–1999)

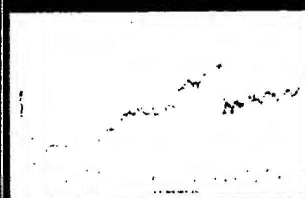

Black Monday of October 1987
(Dow Jones Industrial 1985–1990)

Japanese Stock Market Crash in 1990
(Nikkei 225 1981–1999)

Mark the Position to Future

In the absence of history with respect to Internet stocks, proxy scenarios are forecast from comparable past events to cover a wide spectrum of future outcomes.

The three-year period following each historic market peak serves as the proxy for downside risk while the one-year period that precedes each historic market peak serves as the proxy for further growth scenarios.

Downside scenarios are bootstrapped using daily returns generated from the downside proxies, upside scenarios are bootstrapped using daily returns generated from the upside proxies.

Scenario	Day 1	Day 10	Day 11	Day 30	Day 31	Day 60	Day 61	Day 99	Day 100
Scenario 4	[illegible]	[illegible]		[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Scenario [illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Nokia A	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	
Scenario 6	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]		[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Scenario 171	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Nikkei 225 Crash (1980)	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Scenario 91	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Scenario [illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Gold Crash 1980	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]		[illegible]	
Scenario [illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]		[illegible]	[illegible]	[illegible]
Scenario [illegible]		[illegible]			[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Black Monday 1987 [illegible]				[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Scenario [illegible]		[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Scenario [illegible]	[illegible]		[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	
Scenario [illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Stock Market Crash of 1929	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Scenario [illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
Scenario [illegible]	[illegible]	[illegible]	[illegible]	34.4[illegible]	96.36	[illegible]	[illegible]	[illegible]	[illegible]
Scenario 184	[illegible]	[illegible]	[illegible]	35.86	95.37	[illegible]	[illegible]	22.03	[illegible]
Scenario 179	146.29	[illegible]	136.22	96.62	98.95	53.16	52.43	20.66	20.48

and volatility scenarios are bootstrapped by combining returns from the upside and downside proxies. This process creates a set of 200 scenarios over a horizon of 100 days.

Assuming a hold strategy, the MtF values of one share of Amazon.com indicate that the stock can potentially gain 120% in the next 100 days. However, it can also lose more than 85% over the same period. This translates into a loss of more than $587 million in just five months on a position of five million shares! When sorted in decreasing order according to the magnitude of the gain (or loss), the potential upside and downside risk is apparent.

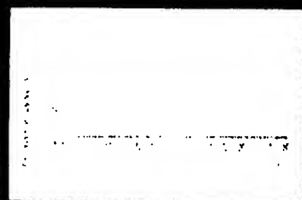

Scenario 10 Days Mark-to-Future

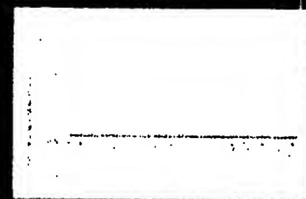

Scenario 30 Days Mark-to-Future

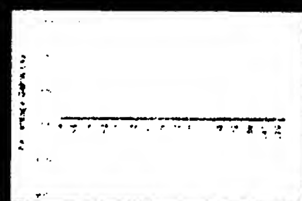

Scenario 60 Days Mark-to-Future

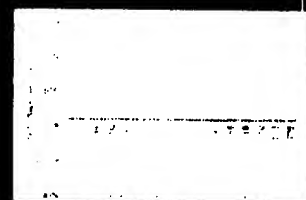

Scenario 100 Days Mark-to-Future

Would a trading strategy mitigate risk? Five million shares represent 10% of Amazon.com's total market float. The average daily volume for Amazon.com is only seven million shares, it is common for daily trading volume to be as low as 10% of average volume. This lack of liquidity poses additional concerns–instantaneous disposition of the entire holdings is not likely. A stress test model must forecast the risk specific to the Amazon.com shares, as well as the risks of trading in this illiquid market.

To account for a light trading day in which daily volume is limited to 700,000 shares, the strategy is conditioned on the movement of the Amazon.com share price. Whenever Amazon.com falls below 10% of its current value, the position size is decreased by 250,000 shares. When the price drops below 20% of its current value, the disposition of shares is increased to 500,000. For simplicity, proceeds from the sale remain as uninvested cash.

Two bid/ask spread curves depict normal and extreme liquidity risks. The first curve portrays limited liquidity risks by assuming that sales can be transacted without moving market prices by 20 more than basis points. The second curve portrays a scenario of high liquidity risk corresponding to a discount of 20% to the market price of Amazon.com for a trading volume of 700,000 shares.

The dynamic trading strategy yields benefits. Under the worst downside scenario, the hold strategy forecasts a potential loss of 85% over the next 100 days. By locking in some of the gains, potential losses are limited to 28% under normal liquidity conditions and 38% in a highly illiquid market. ⊙

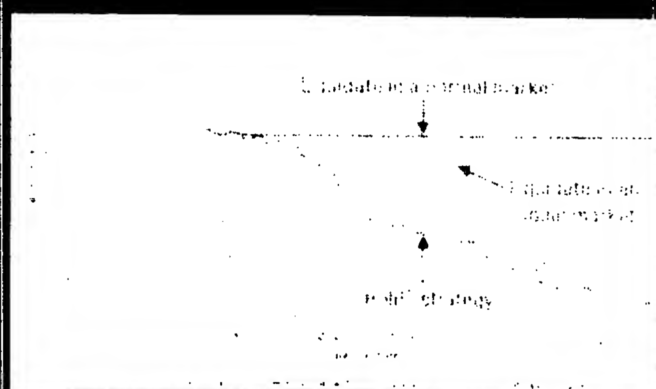

Impact of Liquidity (the next 100 days)

Bootstrapping and Proxies

If risk measures based on a summary statistic such as standard deviation would not have captured the exchange risk in the Canadian dollar in March 1998, what methods might have been used to determine the six-month risk in the CAD/USD exchange rate? Novel scenarios on future Canadian dollar rates for August and September 1998 could have been extracted from an AUD/USD time series. Using the AUD/USD exchange rate as a proxy for the scenarios on the CAD/USD rate could have been justified on the basis of perceived similarities between the Australian and Canadian macroeconomic conditions at that time. Several times over the last decade, for example in 1987, 1989 and 1993, the exchange rate on the Australian dollar dropped by 10% or more within a short number of months. Scenarios on the AUD/USD exchange rate would have provided a forward-looking view of future Canadian dollar behaviour, including a Canadian dollar scenario describing a drop to 0.633 USD or below.

Consider Brazil as another example. Early in 1999, market risk managers in Brazil asked "How can we calculate Value-at-Risk figures for our portfolios when we don't have any historic data?" (Locke 1999). One possibility would have been to consider the market turbulence caused by the devaluation of the *ruble* in August 1998 as a proxy for how events in Brazil might unfold should the *real* be allowed to float (Yung 1999a).

Another way to produce forward-looking scenarios is to use statistical bootstrapping to extend or *fill-in* history. This involves taking random samples from a time series in order to create synthetic time series. For example, a synthetic high-volatility time series can be created by sampling repeatedly from a filtered history that excludes low volatility periods.

Previous markets that have enjoyed exceptional growth, for example, Nokia A stock (an emerging industry) and the gold rush in the 1970s, can serve as proxies for upside scenarios to determine the risk in a position in a new Internet stock (Yung 1999b). The 1929 stock market crash that initiated the Great Depression and Black Monday of October 1987 can serve as proxies for downside scenarios. Bootstrapping is achieved by mixing, scrambling and rearranging daily returns from the applicable proxy periods. Further randomization can be achieved by randomly selecting the starting point for a scenario series.

Thus, by taking a broad perspective on the sources of scenarios, it is possible to extract reasonable forward-looking scenarios directly from history or to use a statistical technique such as bootstrapping.

Implementation Issues

The more history that is used, the more likely the scenarios are to span the range of future events. However, more scenarios mean more data storage and MtF computations. Balancing computational effort with coverage is one of the challenges of scenario generation. Scenario banding is a promising method for reducing the number of historical scenarios required to calculate accurate risk measures (Cartolano and Verma 1999). Further research is required to identify methods to trade-off coverage and computational effectiveness.

Extensive research has been undertaken, with significant success, on reducing the number of scenarios needed for accurate Monte Carlo results. Examples include low discrepancy sequences, applied by Boyle (1977) to option pricing and subsequently applied by Kreinin et al. (1998a) to risk measurement, the application of grids (Chishti 1999) and stratified sampling (Jamshidian and Zhu 1997). Recent research contributions (Kreinin et al. 1998b) extend the scope of these solutions to portfolios with high dimensionality. Practitioners use heuristic methods that extract representative stress scenarios from a much larger scenario set that spans a wide range of potential future outcomes. For example, asset managers routinely map a portfolio into a benchmark index and use the historical outcomes of such a benchmark index to assess and communicate the relative risk of their actual portfolio holdings.

There are two aspects to the forward-looking nature of MtF. The first, the use of forward-looking scenarios, has been discussed in this chapter. In the next chapters we present the second: the simulation that produces the values that populate the MtF Cube. The simulation accounts completely for all the forward events and their effect on the future value of the underlying instruments and for the dynamic portfolio strategies that model changes to the positions in the portfolios.

Produce the Basis MtF Cube

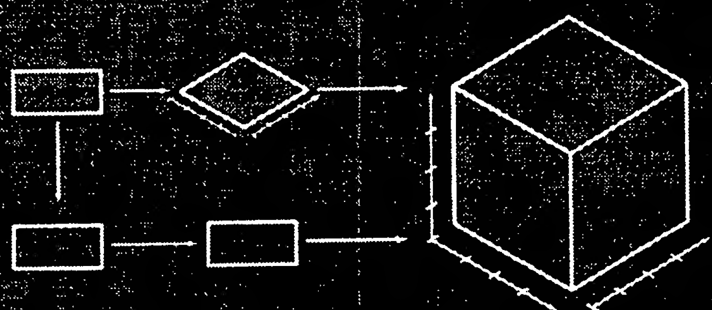

Figure [illegible]: Generation of a Basis MtF Cube

Produce the Product MtF Cube

In order to calculate MtF values for new deals that have been newly transacted or to simply reduce the dimensions of the Basis MtF Cube, it is often advantageous to map actual financial products onto the basis instruments contained in the MtF Cube. The result of this mapping exercise is the creation of a **Product MtF Cube** containing the MtF values of all financial products of interest. Unlike the Basis MtF Cube, the Product MtF Cube does not need to exist in a physical sense as it is completely defined by the product mappings themselves. That is, the cube need not be saved, it only needs to be calculated as required.

It is important to note that the incorporation of this mapping step is simply a configuration choice. It is always possible to define the basis instruments to be actual financial products themselves. The mapping in this step becomes simply one-to-one. In this case, however, the dimensions of the Basis MtF Cube must be of such a magnitude to incorporate all possible financial products of interest.

A product MtF table, M'_t, has dimensions $S_t \times N'$, where S_t is the number of scenarios and N' is the number of financial products. Each cell contains the MtF value, $m'_{i'jt}$, of financial product, i' ($i'=1,\ldots,N'$), under scenario, j ($j=1,\ldots,S_t$) at time t. A product MtF table is constructed as a function of the mapping of each financial product onto the basis MtF table in the following manner:

$$M'_t = f_t\,(M_t)$$

where $f_t\,()$ represents a set of operators mapping the N' financial products onto the N basis instruments at time step t. A Product MtF Cube consists of T product MtF tables corresponding to each time step, t ($t = 1,\ldots,T$).

For many financial products, the mapping is a static function (constant across scenarios and time steps) and a straightforward linear combination of the basis instruments. In this case, the mapping is analogous to static holdings of individual basis instruments in a portfolio. An example is a fixed rate bond that would be mapped simply into a static combination of zero coupon basis instruments associated with the coupon dates of the bond.

For more complex financial products the mapping may be dynamic (varying across scenarios and/or time steps) and a non-linear function of the basis instruments. An example is a floating rate note. Initially it would be mapped into a single zero coupon basis instrument but, at each reset date, it will roll over into a new zero coupon basis instrument in an amount that is a function of the new basis instrument's MtF value under each scenario.

Figure 5.1 illustrates a Product MtF Cube that has been generated by mapping N' financial products onto

a Basis MtF Cube, where $f()$ represents the T sets of operators, f_t, $(f_1, \ldots, f_T)$ associated with the mapping of the N' financial products onto the N basis instruments for each time step t.

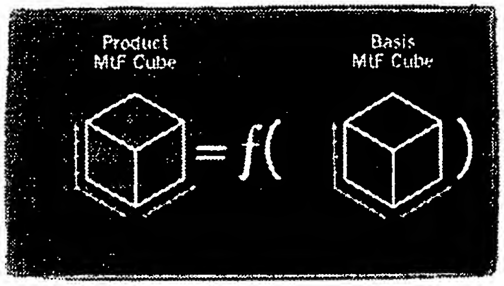

Figure 5.1: Mapping products onto a Basis MtF Cube

When the Basis MtF Cube spans all the relevant risk factors, $(u_1, \ldots, u_k)$, required to generate MtF values for a given financial product, a perfect mapping of the product onto the basis instruments can be achieved. In other cases, when the risk factors are not contained in the MtF Cube, a perfect mapping will not be achieved unless the risk factors are incorporated as part of the mapping exercise. As an example, the implied volatility of a newly transacted OTC option may represent a risk factor not contained in the Basis MtF Cube. Thus, the mapping of this financial product onto the basis instruments requires the incorporation of the implied volatility as an exogenous parameter.

Perfect Mapping of Products onto Basis Instruments

In many cases, the mapping exercise results in a perfect replication of the financial product's MtF values with respect to the values produced by the product's pricing model. This implies that all relevant risk factors associated with the valuation of these products are spanned by the Basis MtF Cube.

Complicating matters somewhat is the fact that even though all risk factors may be contained in the Basis MtF Cube, for some risk factor classes an interpolation amongst the risk factors is required for valuation purposes. Interest rate risk factors represent a typical case. A term structure may contain a number of independent nodes that are associated with the discount factors for given terms. Products with cash flows that coincide precisely with those particular terms can be mapped one-to-one onto zero coupon bond basis instruments whose maturities also fall precisely on those terms.

However, products with cash flows that do not coincide with those particular terms require a slightly more complex mapping. This issue from a mapping perspective is completely analogous to the interpolation issues faced when valuing products that require a term structure as a pricing model input. In this case, the mapping exercise can be broken down into two conceptual components; the intermediate mapping of the financial product onto hypothetical basis instruments followed by the mapping of each hypothetical basis instrument onto actual basis instruments contained in the MtF Cube.

Mapping Financial Products onto Hypothetical Basis Instruments

As an example, the fixed leg of an OTC interest rate swap can be mapped onto a static set of hypothetical zero coupon basis instruments (bundled) whose maturities correspond precisely to the cash flow dates of the swap. Consider a swap with maturity at $t = 3$, notional n, and a fixed rate, r_f, corresponding to a term of $t = 1$. Figure 5.2 illustrates the mapping sequence used to perfectly map this financial product onto the hypothetical basis instruments.

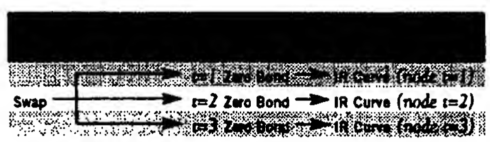

Figure 5.2: Mapping sequence of a swap onto zero coupon basis instruments

The MtF value of the fixed leg of the swap ($i'=1$) can be determined as a static function of the zero coupon basis instruments ($i=1$, 2, and 3) associated with time steps $t=1$, 2, and 3. In this case, the function is a linear combination of the hypothetical basis instruments and can be represented as a function of their MtF values $(m_{1jt}, m_{2jt}, m_{3jt})$:

$$\begin{aligned} m'_{1jt} &= f(m_{1jt}, m_{2jt}, m_{3jt}) \\ &= r_f \cdot n \cdot m_{1jt} + r_f \cdot n \cdot m_{2jt} + (1 + r_f) \cdot n \cdot m_{3jt} \end{aligned}$$

for all time steps t.

Bundled t=1 Zero Bond	$d_{0,1}$	1	$f_{1,2}$	$f_{1,2}f_{2,3}$	$f_{1,2}f_{2,3}f_{3,4}$
Bundled t=2 Zero Bond	$d_{0,2}$	$d_{1,2}$	1	$f_{2,3}$	$f_{2,3}f_{3,4}$
Bundled t=3 Zero Bond	$d_{0,3}$	$d_{1,3}$	$d_{2,3}$	1	$f_{3,4}$
Swap Fixed Leg	$r_F n d_{0,1} + r_F n d_{0,2} + (1+r_F) n d_{0,3}$	$r_F n + r_F n d_{1,2} + (1+r_F) n d_{1,3}$	$r_F n f_{1,2} + r_F n + (1+r_F) n d_{2,3}$	$r_F n f_{1,2} f_{2,3} + r_F n f_{2,3} + (1+r_F) n$	$r_F n f_{1,2} f_{2,3} f_{3,4} + r_F n f_{2,3} f_{3,4} + (1+r_F) n f_{3,4}$

Table 5.1: MtF values of a swap fixed leg across scenarios

($d_{t,t+x}$ and $f_{t-x,t}$ represent the MtF values of the zero coupon bond basis instruments. $d_{t,t+x}$ is the discount factor from time step t to $t+x$ and $f_{t-x,t}$ is the future value factor from time step $t-x$ to t, along a given scenario. The MtF value of the swap fixed leg is the sum of the MtF values of the bundled basis instruments.)

Table 5.1 summarizes the swap fixed leg MtF values derived as a function of the basis instrument MtF values across a given scenario path for the T=4 time horizon.

In contrast to the fixed leg, the floating leg of the same swap must be mapped onto this set of hypothetical zero coupon instruments as a dynamic non-linear function across scenarios and time steps. The floating leg mappings change at each reset date (in this example, at t=1, 2) and are a function of the MtF value of the appropriate zero coupon basis instrument under each scenario. Prior to the first reset date ($t<1$) when the rate has already been preset at r_{-1}, (at time $t=1$), the MtF value of the floating leg ($i'=2$) can be represented using its 'fixed notional' representation, as the linear function

$$
\begin{aligned}
m'_{2jt} &= f(m_{1jt}, m_{2jt}, m_{3jt}) \\
&= (1 + r_{-1}) \cdot n \cdot m_{1jt} \\
&for \quad t < 1
\end{aligned}
$$

From the first reset date ($t=1$) until just prior to the second reset date ($t<2$), the MtF value of the floating leg can be represented as the non-linear function

$$
\begin{aligned}
m'_{2jt} &= f(m_{1jt}, m_{2jt}, m_{3jt}) \\
&= r_{-1} \cdot n \cdot m_{1jt} + \frac{1}{m_{2j1}} \cdot n \cdot m_{2jt} \\
&for \quad 1 \le t < 2
\end{aligned}
$$

The first term in this expression is the value of the reinvested first coupon payment made at t=1, the second term, the value of the remaining swap payments (again using the 'fixed notional' representation). From the second reset date ($t=2$) onward, the MtF value of the floating leg can be represented as the non-linear function

$$
\begin{aligned}
m'_{2jt} &= f(m_{1jt}, m_{2jt}, m_{3jt}) \\
&= r_{-1} \cdot n \cdot m_{1jt} + \left(\frac{1}{m_{2j1}} - 1 \right) \cdot n \cdot m_{2jt} + \frac{1}{m_{3j2}} \cdot n \cdot m_{3jt} \\
&for \quad t \ge 2
\end{aligned}
$$

The first two terms in this expression are the value of the reinvested first and second coupons, respectively, the third term, the value of the final swap payment. Note that the dynamic mapping of the swap floating leg is equivalent to a strategy of rolling over zero coupon basis instruments at each reset date. At any given time step the value of the floating leg is equal to a position in a zero coupon instrument. Until $t=1$ this is a known amount. At the first reset date ($t=1$) the position rolls into a new zero coupon instrument (with maturity at the next reset date t=2), in an amount determined strictly by the MtF value of the zero coupon bond m_{2j1}, under each appropriate scenario. The amount of dynamic rebalancing at each future reset date is always a function of the MtF values of the appropriate zero coupon instrument at those dates.

Table 5.2 illustrates the swap floating leg MtF values derived as a function of the hypothetical basis instrument MtF values across a given scenario path for the T=4 time horizon.

As another example, a foreign exchange (FX), forward can also be mapped onto a static set of domestic and foreign hypothetical zero coupon basis instruments. Consider an FX forward whereby a notional of n in the foreign currency is received at $t = 3$ in exchange for an amount associated with a strike price of k (representing the locked-in forward exchange rate expressed in the domestic currency). Figure 5.3 illustrates the mapping sequence used to replicate this financial product.

Bundled $t=1$ Zero Bond	$d_{0,1}$	1	$f_{1,2}$	$f_{1,2}f_{2,3}$	$f_{1,2}f_{2,3}f_{3,4}$
Bundled $t=2$ Zero Bond	$d_{0,2}$	$d_{1,2}$	1	$f_{2,3}$	$f_{2,3}f_{3,4}$
Bundled $t=3$ Zero Bond	$d_{0,3}$	$d_{1,3}$	$d_{2,3}$	1	$f_{3,4}$
Swap Floating Leg	$(1+r_{.1})nd_{0,1}$	$r_{.1}n + (n/d_{1,2})d_{1,2}$	$r_{.1}n f_{1,2} + n(1/d_{1,2} - 1) + (n/d_{2,3})d_{2,3}$	$r_{.1}n f_{1,2}f_{2,3} + n(1/d_{1,2} - 1) f_{2,3} + n$	$r_{.1}n f_{1,2}f_{2,3}f_{3,4} + n(1/d_{1,2} - 1) f_{2,3}f_{3,4} + n f_{3,4}$

Table 5.2: MtF values of a swap floating leg across scenarios

($d_{t,t+x}$ and $f_{t-x,t}$ represent the MtF values of the zero coupon bond basis instruments. $d_{t,t+x}$ is the discount factor from time step t to $t+x$ and $f_{t-x,t}$ is the future value factor from time step $t-x$ to t, along a given scenario. The MtF value of the swap floating leg is the sum of the MtF values of the bundled basis instruments.)

Bundled $t=3$ Foreign Bond	$q_0 d^*_{0,3}$	$q_1 d^*_{1,3}$	$q_2 d^*_{2,3}$	q_3	$q_4 f^*_{3,4}$
Bundled $t=3$ Domestic Bond	$d_{0,3}$	$d_{1,3}$	$d_{2,3}$	1	$f_{3,4}$
FX Forward	$nq_0 d^*_{0,3} - knd_{0,3}$	$nq_1 d^*_{1,3} - knd_{1,3}$	$nq_2 d^*_{2,3} - knd_{2,3}$	$nq_3 - kn$	$nq_4 f^*_{3,4} - knf_{3,4}$

Table 5.3: MtF values of an FX forward across scenarios

(where $d^*_{t,t+x}$ is the foreign discount factor from time step t to $t+x$ and where q_t represents the spot FX rate (expressed in the domestic currency) at observed at t)

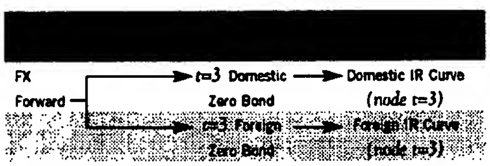

Figure 5.3: Mapping sequence of FX forward onto zero coupon basis instruments

The MtF value of the FX forward ($i' = 3$) can be determined as a static function of the domestic and the foreign zero coupon basis instruments ($i = 3$ and 4) associated with time step $t = 3$. In this case the function is a linear combination of the hypothetical basis instruments and can be represented as follows:

$$\begin{aligned} m'_{3jt} &= f(m_{3jt}, m_{4jt}) \\ &= n \cdot m_{4jt} - k \cdot n \cdot m_{3jt} \end{aligned}$$

for all time steps t.

Table 5.3 illustrates the FX forward MtF values derived as a function of the hypothetical basis instrument MtF values across a given scenario path for the $T = 4$ time horizon.

The straightforward mappings of swaps and FX forwards onto zero coupon basis instruments presented above is possible because basis instruments exist with maturity dates corresponding to all required cash flow dates. When a hypothetical basis instrument coincides with an actual basis instrument in the MtF Cube, the mapping need not be more complex than the cash flow examples presented above.

Mapping of Hypothetical Basis Instruments onto Actual Basis Instruments

However, if no actual basis instruments coincide with the hypothetical basis instruments, a second component to the mapping exercise is required. This additional component maps the hypothetical basis instrument onto the actual basis instruments contained in the Basis MtF Cube. The choice of the mapping function for interest rate basis instruments is equivalent to the choice of interpolation between two term structure nodes.

Consider a one-unit hypothetical zero coupon basis instrument with a maturity date $2<t^*<3$ that must be mapped onto two actual zero coupon basis instruments with maturities of $t=2$ and $t=3$, respectively. This second component of the mapping exercise is illustrated in Figure 5.4.

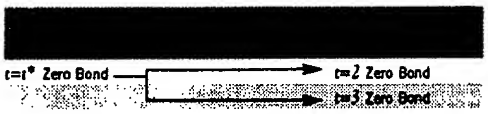

Figure 5.4: Mapping of hypothetical basis instruments onto actual basis instruments

One possible mapping function is based on linear interpolation of a term structure at the level of the discount factor. The MtF value of the hypothetical basis instrument ($i=i^*$), in this case, can be determined as a static linear function of the actual basis instruments ($i=2$ and 3) and can be represented as follows:

$$
\begin{aligned}
m_{i^*jt} &= f(m_{2jt}, m_{3jt}) \\
&= \left(\frac{t_3-t^*}{t_3-t_2}\right)\cdot m_{2jt} + \left(\frac{t^*-t_2}{t_3-t_2}\right)\cdot m_{3jt}
\end{aligned}
$$

for all time steps t.

A second possible mapping function is identical to the linear interpolation of a term structure at the level of the continuously compounded rate of return. The MtF value of the hypothetical basis instrument ($i=i^*$), in this second case, can be determined as a static non-linear function of the actual basis instruments ($i=2$ and 3) and can be represented as follows:

$$
\begin{aligned}
m_{i^*jt} &= f(m_{2jt}, m_{3jt}) \\
&= m_{2jt}^{\left(\frac{t_3-t^*}{t_3-t_2}\right)\frac{t^*}{t_2}} \cdot m_{3jt}^{\left(\frac{t^*-t_2}{t_3-t_2}\right)\frac{t^*}{t_3}}
\end{aligned}
$$

for all time steps t.

Including the second component of the mapping significantly reduces the number of instruments required in the Basis MtF Cube. Perfect mapping can still be achieved if a zero coupon basis instrument exists with maturities corresponding to each independent node of the term structure. As all independent interest rate risk factors are contained in the Basis MtF Cube, the number of basis instruments never needs to exceed the number of nodes in the term structure.

As an example, the interest rate risk factors contained in the RiskMetrics data files imply a set of zero coupon basis instruments for government and interbank risk of selected currencies. Any arbitrary cash flow can, therefore, be mapped onto hypothetical basis instruments which are subsequently mapped onto the RiskMetrics implied actual basis instruments in a fully consistent manner.

Note that each of the product mappings described in Perfect Mapping of Products onto Basis Instruments produces perfect MtF value replications for the financial products across all scenarios and time steps.

Significantly, the mapping of the swap floating leg fully captures the reset risk inherent in the product as it is effectively replicated by a dynamic rollover strategy of the basis instruments. In these examples, all risk inherent in these products is embodied in the basis instruments that are contained in the Basis MtF Cube. As a consequence, the valuation of these OTC securities (which may not be known prior to generating the Basis MtF Cube) are defined strictly by their mappings onto the basis instruments. They do not need to be simulated individually.

Imperfect Mapping of Products onto Basis Instruments

In certain cases, the mapping of financial products onto basis instruments results in imperfect replication. This may be because a less complex mapping function has been chosen for a specific application. On the other hand, it may be simply because the Basis MtF Cube does not contain basis instruments associated with all relevant risk factors.

In the latter case, the MtF value of a given financial product must be determined as a function of the pre-computed MtF values of the basis instruments as well as a set of additional risk factors, $u^*_{k^*} = u_1, \ldots, u^*_{K^*}$, which are not directly contained in the MtF Cube:

$$m'_{ijt} = f(m_{1jt}, \ldots, m_{Njt}, u_1^*, \ldots, u^*_{k^*})$$

In this case, the mapping exercise itself must incorporate the additional risk factors, either statically or dynamically, across scenarios and time steps.

Consider a newly transacted OTC option whose implied volatility represents a risk factor that is not associated with a basis instrument contained in the Basis MtF Cube. To produce a correct MtF value for this product, an additional volatility risk factor $u^*_{k^*}$ must be incorporated into the product mapping function itself. Note that the volatility risk factor may be assumed constant ($u^*_{k^*}$ is constant across scenarios and time steps) or assumed to be associated with a 'smile' or a 'skew'. In this case, $u^*_{k^*} = g(M_t)$ becomes a function of MtF values across scenarios and time steps.

As an example, the mapping of an FX option requires two inputs: an implied volatility (associated with the foreign currency) and mappings onto a set of domestic and foreign zero coupon basis instruments. Consider a European FX call option with the same notional n, maturity $t=3$ and strike price k as that of the FX forward described in the previous section. Figure 5.5 illustrates

the mapping sequence used to map the option onto the basis instruments in the MtF Cube.

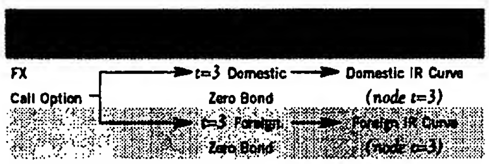

Figure 5.5: Mapping sequence of FX call option onto zero coupon basis instruments

The MtF value of the FX option ($i'=4$) can be determined as a function of the implied volatility along with the domestic and the foreign zero coupon basis instruments ($i=3$ and 4) associated with time step $t=3$. In contrast to the FX forward, the mapping of the FX option requires a non-linear function that can be based on the Garman—Kohlhagen model for currency derivatives (Garman and Kohlhagen 1983):

$$
\begin{aligned}
m'_{4jt} &= f(m_{3jt}, m_{4jt}, u_1^*) \\
&= n \cdot cdf\left(\frac{\ln\left(\frac{m_{4jt}}{k \cdot m_{3jt}}\right) + \frac{u_1^{*2} \cdot (3-t)}{2}}{u_1^* \sqrt{3-t}} \right) \cdot m_{4jt} \\
&\quad - k \cdot n \cdot cdf\left(\frac{\ln\left(\frac{m_{4jt}}{k \cdot m_{3jt}}\right) - \frac{u_1^{*2} \cdot (3-t)}{2}}{u_1^* \sqrt{3-t}} \right) \cdot m_{3jt}
\end{aligned}
$$

where u^*_1 represents the exogenous implied volatility risk factor associated with the foreign currency and $cdf(\ )$ is the normal cumulative distribution function. Note that although volatility is assumed constant in this case, the mapping fully incorporates the non-linear relationship between this option and the underlying risk factors.

This mapping can be extended to incorporate a volatility 'smile' by including the exogenous risk factor as a function of the MtF values of the basis instruments or

$$u_1^* = f^*(m_{3jt}, m_{4jt})$$

While the above mapping produces an upper bound with respect to pricing accuracy, a high degree of complexity is placed upon the mapping step itself. Other mapping approaches may trade-off full pricing accuracy for less mapping complexity. A straightforward example is to map the option onto the same basis instruments described above, but based instead upon a *non-dynamic* delta (or a delta that is adjusted only for certain time steps).

A second example is to generate a series of abstract call options (basis instruments) associated with a generic underlying security (normalized to a MtF value of one) and corresponding to varying volatilities and terms to expiry. The mapping function, in this case, may be captured by a linear combination of the abstract basis instruments.

This chapter has illustrated the benefits and some of the challenges associated with mapping financial products onto the Basis MtF Cube. The key benefit is the ability to significantly reduce the dimensionality of the Basis MtF Cube and to provide a mechanism for incorporating newly transacted OTC instruments.

Consider a financial institution with an inventory of 100,000 vanilla swaps (of a common currency) and the desire to produce a risk report by simulating across 1,000 scenarios and over 100 time steps. Without the incorporation of a product mapping step, the dimensions of the required MtF Cube must be 100,000 instruments x 1,000 scenarios x 10 time steps. The number of cells contained in the cube and the number of re-valuations required will be one billion.

The inclusion of a step mapping financial products onto basis instruments can reduce the magnitude of this exercise significantly. If we assume that the institution utilizes a 10-node term structure (and linearly interpolates for other term nodes) in the pricing of these swaps it implies that the number of basis instruments can be reduced to ten without loss of accuracy. The dimensions of the required Basis MtF Cube are, thus, reduced to 10 instruments x 1,000 scenarios x 10 time steps with number of re-valuations reduced dramatically to 100,000, a decrease in the computational effort by a factor of 10,000.

The key challenges of mapping financial products onto the Basis MtF Cube are associated with developing appropriate mapping functions that optimally trade off accuracy and complexity for specific applications. The mapping of financial products onto basis instruments occurs strictly in the post-Cube stage of the MtF framework, thus enabling many possible mapping functions to be developed and incorporated in an application. Clearly, this is an area where there is great scope for further development. ⊛

Produce the Portfolio MtF Cube

In order to determine the MtF values of a static portfolio or a portfolio regime through time, the next step in the MtF framework involves the mapping of portfolio strategies onto the Product MtF Cube. The result of this second mapping exercise is the creation of a **Portfolio MtF Cube** containing the MtF values of all portfolio regimes of interest. Similar to the Product MtF Cube, the Portfolio MtF Cube need not exist in a physical sense as it is completely defined by the combined portfolio and product mappings.

A portfolio MtF table corresponding to time t, M''_t, has dimensions $S_t \times N''$, where S_t is the number of scenarios and N'' is the number of portfolio regimes. Each cell contains the MtF value, $m''_{i''jt}$, of portfolio regime i'' ($i''=1,\ldots, N''$), under scenario j ($j=1,\ldots, S_t$). Each portfolio MtF table is generated as a function of the product MtF table and an $N' \times N''$ matrix of portfolio positions, x_t, with each cell $x_{i'i''t}$ containing the mapping of portfolio regime i'' onto financial product i' at time step t. A Portfolio MtF Cube consists of T portfolio MtF tables corresponding to each time step t ($t=1,\ldots,T$).

Only for a static *buy-and-hold* portfolio regime are the positions, x_t, constant for all t. In general, given that time is an explicit dimension of the MtF framework, position size may be time dependent, thus enabling a *dynamically* changing portfolio regime to be modeled. The fifth step of the MtF framework incorporates the mapping of portfolio regimes onto the financial products as indicated by the mapping sequence in Figure 6.1.

Portfolio Regimes $\longrightarrow$ **Financial Products**

Figure 6.1: Portfolio mapping sequence

Figure 6.2 illustrates a Portfolio MtF Cube comprising the MtF values for individual portfolios generated by mapping onto the MtF values of the financial products at each time step.

Note that a given Portfolio MtF Cube may contain several different regimes that correspond to the same initial portfolio at $t=0$ (i.e. $x_{i'i''0}$ may be the same for several different regimes) as illustrated in Figure 6.3.

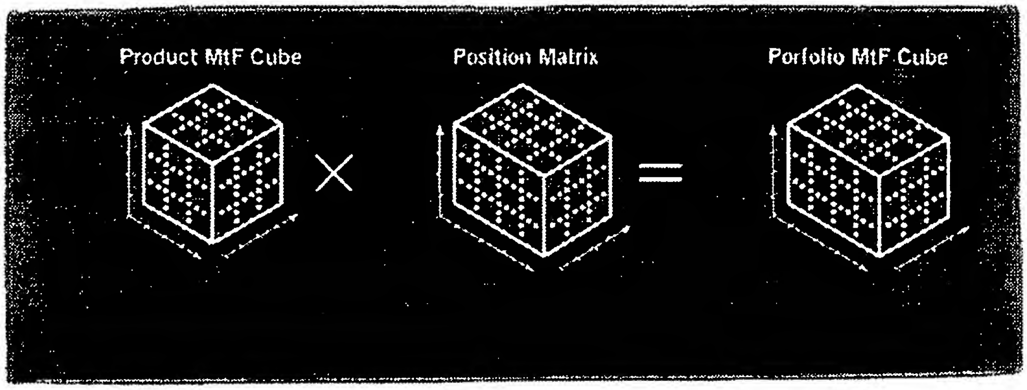

Figure 6.2: Mapping of Portfolio MtF Cube onto a Product MtF Cube

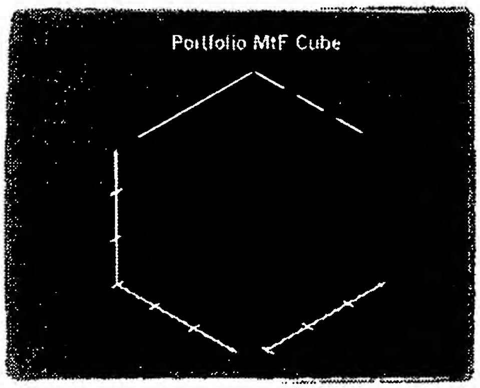

Figure 6.3: Representation of multiple portfolio regimes

In general, the mapping of portfolio positions can be categorized into two general strategies, predetermined regimes and conditional regimes.

Predetermined Regimes

In a **predetermined regime**, portfolio strategies are independent of the contents of the Basis (or Product) MtF Cube; the quantity of positions to be mapped onto each financial product at each time step is defined by a simple position schedule. A portfolio MtF table can, therefore, be constructed as a straightforward linear combination of a fixed quantity of financial products in the product MtF table:

$$M''_t = M'_t x_t$$

where, each cell of the portfolio positions matrix, $x_{i'i''t}$, contains the number of positions of product i' held in portfolio i''.

Note again that a buy-and-hold strategy is simply a special case of a predetermined regime where the portfolio positions remain constant for all t (i.e., $x_{i'i''t}$ is constant for all t).

As a more complex example, consider the modeling of liquidity risk. If the passage of time or dynamic portfolio mapping cannot be modeled in the risk assessment framework, the change in mark-to-market value or a portfolio's unrealized P&L serves as a proxy for its realized P&L or the actual liquidation of portfolio positions. A liquidity-adjusted risk measure is often estimated by applying a simplistic add-on to a non-adjusted risk measure. In contrast, the MtF framework allows for the explicit liquidation of holdings through specification of a regime that liquidates the portfolio holding x_t over a given time horizon.

Consider a single bond ($i'=1$) with a MtF value of m'_{1jk}. Different portfolio liquidation assumptions can be assessed by applying different regimes. The portfolio mapping sequence associated with a portfolio regime is illustrated in Figure 6.4.

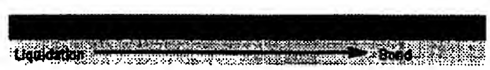

Figure 6.4: Mapping of portfolio liquidation regime onto a bond

Under the first portfolio regime ($i''=1$), the bond is assumed to be liquid and can be fully liquidated by $t=1$. In this case, the MtF bond position value at $t=0$ is simply

$$\begin{aligned} m''_{1j0} &= x_{110} \cdot m'_{1j0} \\ &= 1 \cdot m'_{1j0} \\ &= m'_{1j0} \\ for \quad t &= 0 \end{aligned}$$

and after full liquidation at $t \geq 1$

$$\begin{aligned} m''_{1jt} &= x_{11t} \cdot m'_{1jt} \\ &= 0 \cdot m'_{1jt} \\ &= 0 \\ for \quad t &\geq 1 \end{aligned}$$

Under the second portfolio regime ($i''=2$), the bond is assumed to be illiquid and can only be liquidated by 10% over each $t=1,\ldots,10$. In this case, the MtF bond position value from $0 \leq t \leq 10$ can be represented as

$$\begin{aligned} m''_{2jt} &= x_{21t} \cdot m'_{1jt} \\ &= (1-0.1t) \cdot x_{210} \cdot m'_{1jt} \\ for \quad 0 &\leq t \leq 10 \end{aligned}$$

and after full liquidation at $t > 10$

$$\begin{aligned} m''_{1jt} &= x_{21t} \cdot m'_{1jt} \\ &= 0 \cdot m'_{1jt} \\ &= 0 \\ for \quad t &> 10 \end{aligned}$$

Note that for a common time horizon of $T=10$, products with differing liquidation periods may now be assessed within a consistent framework through the dynamic mapping of portfolios onto financial products. The calculation of a particular liquidation risk measure will typically require the settlement into a cash account in addition to the mechanism of bond liquidation described in this example. Settlement into a cash account is dependent upon the bond MtF values under each scenario and, thus, represents an example of a conditional strategy as described in the next section.

As a second example of a pre-determined regime, consider an attribution or backtesting analysis for a portfolio over a given time period. Under a standard framework, the portfolio's P&L impact due to market risk factor changes over this time period may be captured solely through the appropriate selection of scenarios (for the current and the previous period market risk factor levels). In addition to these factors, however, the MtF framework enables the capture of the P&L impact due to time decay and position change by the explicit incorporation of time and different portfolio regimes, respectively.

In an **attribution analysis** the impact of position change on a portfolio from the previous day ($t=0$) until today ($t=1$) is assessed. The analysis consists of two scenarios, the previous day's risk factor levels ($j=0$) and today's risk factor levels ($j=1$), applied to a portfolio that contains a single financial product ($i'=1$) with MtF values of m'_{1jt}.

An attribution or backtesting analysis can be performed by assessing the individual portfolio MtF values that result from the mapping of two position schedules associated with the change in position from $t=0$ until $t=1$. The first portfolio regime ($i''=1$) consists of a static portfolio based strictly on the previous day's positions, or $x_{110} = x_{111}$. The second portfolio regime ($i''=2$) consists of a static portfolio based strictly on today's positions, or $x_{210} = x_{211}$. Table 6.1 contains the portfolio MtF values appropriate for an attribution or backtesting analysis.

Market Factors	m''_{110}	$x_{110} \cdot m'_{110}$
Time Decay	m''_{101}	$x_{111} \cdot m'_{101}$
Position Change	m''_{200}	$x_{210} \cdot m'_{100}$
Total	m''_{211}	$x_{211} \cdot m'_{111}$

Table 6.1: Portfolio MtF values for an attribution analysis

Conditional Regimes

In a **conditional regime** the portfolio strategies depend on the contents of the MtF Cube. In these portfolio strategies, the quantity of positions to be mapped onto each financial product at each time step may be a function of the MtF values in the Basis, Product or Portfolio MtF Cube under given scenarios and time steps. A portfolio position matrix in this case can be represented as a function of the Product MtF Cube in the following manner:

$$x_t = h_t(M')$$

where $h_t()$ represents a set of **re-balancing operators** mapping the portfolio position matrix onto the Product MtF Cube at time step t.

Consider a simple T-Bill rollover strategy. At the $t=1$ maturity of a current T-bill (with an original term of two periods), the proceeds of notional N are reinvested in a second T-bill that matures at $t=3$. At the $t=3$ maturity of the second T-bill, the proceeds are reinvested into a third T-bill maturing at $t=5$. While the rollover strategy is known at each rollover date, the actual quantities of the three T-bills ($i' = 1, 2, 3$) are a function of their MtF values (m'_{1jt}, m'_{2jt}, m'_{3jt}). The portfolio mapping sequence is illustrated in Figure 6.5.

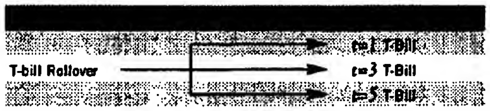

Figure 6.5: Mapping of portfolio rollover strategy onto T-bills

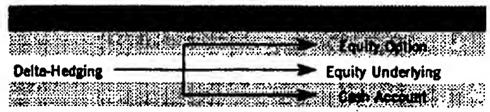

Figure 6.6: Mapping of portfolio delta-hedging strategy onto underlying and cash

Prior to the first rollover date ($t<1$), the MtF value of the portfolio ($i''=1$) is the simple linear function

$$\begin{aligned} m''_{1jt} &= x_{11t} \cdot m'_{1jt} \\ &= 1 \cdot m'_{1jt} \\ &= m'_{1jt} \\ for \quad & t<1 \end{aligned}$$

From the first rollover date ($t=1$) until just prior to the second rollover date ($t<3$), the MtF value of the dynamically rebalanced portfolio can be represented as the non-linear function

$$\begin{aligned} m''_{1jt} &= x_{12t} \cdot m'_{2jt} \\ &= \frac{N}{m'_{2j1}} \cdot m'_{2jt} \\ for \quad & 1 \le t < 3 \end{aligned}$$

From the second rollover date ($t=3$) onward, the MtF value of the portfolio can be represented as the non-linear function

$$\begin{aligned} m''_{1jt} &= x_{13t} \cdot m'_{3jt} \\ &= \frac{N^2}{m'_{2j1} \cdot m'_{3j3}} \cdot m'_{3jt} \\ for \quad & t \ge 3 \end{aligned}$$

The second equation illustrates that at the first reset date ($t=1$) the position rolls into a new T-bill (with maturity at the next rollover date $t=3$) in an amount determined by the MtF value of the T-bill, m'_{2j1}, under each appropriate scenario. The second and third equations illustrate that the amount of dynamic rebalancing at each rollover date depends on the MtF values of the appropriate T-bill at those dates.

Consider as a second example a delta-hedging regime applied to a portfolio consisting of an equity option ($i'=1$), the underlying equity ($i'=2$) and a cash account ($i'=3$). At $t=0$ the portfolio contains a single position in the option and a position in the underlying equity such that the overall portfolio is delta neutral. The cash account at $t=0$ has a position of zero. According to the conditional delta-hedging strategy at $t=2$ and $t=4$ the portfolio is dynamically rebalanced so as to delta-hedge the option. This is achieved by acquiring a position in the equity (funded by the cash account) that returns the overall portfolio to a delta-neutral position. The portfolio mapping sequence for this example is illustrated in Figure 6.6.

To execute a delta-hedging strategy the MtF Cube must contain other MtF measures in addition to MtF values. Specifically, the MtF deltas for each of the financial products ($i'=1, 2, 3$) are required. These are defined as $\Delta'_{i'jt} = \Delta'_{1jt}, \Delta'_{2jt}, \Delta'_{3jt}$. Prior to the first rebalancing date ($t=2$), the MtF value of the portfolio ($i''=1$) can be represented as

$$\begin{aligned} m''_{1jt} &= x_{11t} \cdot m'_{1jt} + x_{12t} \cdot m'_{2jt} \\ &= m'_{1jt} - \Delta'_{1j0} \cdot m'_{2jt} \\ for \quad & t \le 2 \end{aligned}$$

From the first rebalancing date ($t=2$) until just prior to the second rebalancing date ($t<4$), the MtF value of the dynamically re-balanced portfolio can be represented as

$$\begin{aligned} m''_{1jt} &= x_{11t} \cdot m'_{1jt} + x_{12t} \cdot m'_{2jt} + x_{13t} \cdot m'_{3jt} \\ &= m'_{1jt} - \Delta'_{1j2} \cdot m'_{2jt} + \left(\Delta'_{1j2} - \Delta'_{1j0}\right) \cdot \left(\frac{m'_{2j2}}{m'_{3j2}}\right) \cdot m'_{3j2} \\ for \quad & 2 \le t < 4 \end{aligned}$$

From the second rebalancing date ($t=4$) onward, the MtF value of the dynamically rebalanced portfolio can be represented as

$$\begin{aligned} m''_{1jt} &= x_{11t} \cdot m'_{1jt} + x_{12t} \cdot m'_{2jt} + x_{13t} \cdot m'_{3jt} \\ &= m'_{1jt} - \Delta'_{1j4} \cdot m'_{2jt} + \left(\left(\Delta'_{1j2} - \Delta'_{1j0}\right) \cdot \left(\frac{m'_{2j2}}{m'_{3j2}}\right) + \left(\Delta'_{1j4} - \Delta'_{1j2}\right) \cdot \left(\frac{m'_{2j4}}{m'_{3j4}}\right)\right) \cdot m'_{3jt} \\ for \quad & t \ge 4 \end{aligned}$$

The portfolio MtF values in this example not only incorporate the impact of a conditional delta hedging strategy but, additionally, enable the assessment of funding risk by inspection of the MtF values of the cash account. The conditional strategy has been mapped into the scenarios that underlie the MtF Cube and has thus satisfied the requirement that it be a deterministic function of scenarios. A complex multi-stage, multi-period stochastic problem has become a single stage, multi-period problem that can be readily analyzed. The effectiveness of different trading regimes, such as delta and delta-gamma hedging strategies, can be easily compared. ⊜

Produce the Desired Risk-Reward Measures

This chapter describes how risk-reward measures, as well as other quantitative analytics, may be applied to the Product or Portfolio MtF Cubes. These post-processing analytics are typically embedded in lightweight task-oriented user applications. Two general categories of analytics that may be applied in Step 6 include transforming MtF values into other measures, and calculating descriptors of portfolio MtF distributions. For ease of exposition, in this chapter the notation is relaxed and m_{ijt} will be used to represent the MtF value of the product or portfolio i under scenario j at time step t.

Transforming MtF Values into Other Measures

In certain applications, it is necessary to *transform* the product or portfolio MtF values prior to their input to a specific analysis. The transformation can be any arbitrary function (linear or non-linear) of the portfolio MtF value.

Consider a portfolio with a single position in a financial product that is subject to credit risk. A function that takes the maximum of the portfolio MtF value and zero transforms the MtF value to a measure of **counterparty credit exposure**. In addition, if the portfolio is composed of multiple positions, the transformation itself may also account for netting rules and other applicable credit migration provisions. Note, that for the measurement of counterparty credit exposure, the reinvestment of cashflows is not desired and, thus, the basis instruments are not *bundled* to account for total return as described in step one.

Table 7.1 summarizes the transformation of portfolio MtF values required to provide some example of credit exposure measures. Note that these transformation functions require only the information residing in the pre-computed MtF Cube. For the measure **actual exposure**, a simple transformation is applied to the portfolio MtF value under each scenario and time step. In the case of **potential exposure**, the transformation is slightly more complex; it is based on all portfolio MtF values over the $t \le t^* \le T$ time horizon for a given scenario.

Actual Exposure	$AE_{ijt} = \max\lfloor m_{ijt}, 0\rfloor$
Potential Exposure	$PE_{ijt} = \max_{t<t^*\le T}\lfloor m_{ijt^*} - AE_{ijt}, 0\rfloor$
Total Exposure	$TE_{ijt} = AE_{ijt} + PE_{ijt}$

Table 7.1: Transformation of MtF values to credit exposure measures. (where m_{ijt} is the MtF value of single position portfolio i under scenario j at time step t. t^* represents a time step that is in the interval $t \le t^* \le T$ where T is the time horizon.)

Additional information not included in the Portfolio MtF Cube may be required for more complex transformations. This information can be incorporated, statically or dynamically, in the transformation.

Calculating Distribution Descriptors

Typical post-processing applications calculate statistics characterizing the market risk, credit risk, liquidity risk and performance profiles associated with particular products or particular portfolio regimes. A variety of summary statistics may be calculated by aggregating the MtF values across the scenario and/or time step dimensions of the appropriate MtF Cube.

In many cases, the applications will require the inclusion of scenario weightings as defined by the $S_t \times 1$ probability vector p, where each cell p_j ($p_j = p_1, \ldots, p_S$) contains the probability associated with each state j. It is only in this step of the MtF methodology that probabilities must be explicitly incorporated into the process.

In other cases, additional information such as counterparty default probabilities and recovery rates may also be required. These additional information requirements must be incorporated directly into the algorithms embedded in the post-processing application itself.

The risk and reward measures that may be incorporated in this step may span all dimensions of the Portfolio MtF Cube. Thus, in addition to standard measures which often rely on backward-looking scenarios and ignore the passage of time, the MtF framework enables the incorporation of forward-looking measures that have full access to future distributions of MtF values.

An approach such as this is ideal for the calculation of measures that are based upon relative performance with respect to a benchmark. Consider the measure of **expected downside**, calculated as the expected *under-performance* of a product or a portfolio with respect to the benchmark's performance across a range of scenarios. If we let the functions $(x)^+$ and $(x)^-$ equal $max(x,0)$ and $|min(x,0)|$, respectively, then the value of expected downside at time t for a product or portfolio i (with respect to a benchmark) is equal to,

$$E[D_{it}] = \sum_{j=1}^{S_t} p_j (m_{ijt} - \tau_{jt})^-$$

where τ_{jt} represents the value of an arbitrary benchmark under scenario j at time t. This measure is ideal in the context of evaluating the performance of mutual funds with respect to target payoffs. This measure may be modified to include a coefficient of risk aversion λ to provide the risk-adjusted measure of **regret**, or simply

$$E[D_{it}]^* = \lambda \cdot \sum_{j=1}^{S_t} p_j (m_{ijt} - \tau_{jt})^-$$

A key construct underlying the calculation of the distribution descriptors associated with various risk-reward measures is that the input for each is the same MtF Cube. Thus, the MtF methodology provides a unifying framework for the integration of market risk, credit risk and liquidity risk measures. Selected risk measures associated with each of these risk classes are summarized in Table 7.2.

The forward-looking Portfolio MtF Cube can be used for the assessment of reward and performance as well as for risk measurement. Table 7.3 provides selected reward and performance measures applied as post-processing applications.

The Put/Call Efficient Frontier performance measure defined in Table 7.3 provides a novel means of assessing the trade-off between risk and reward realizable only in a MtF framework. This measure essentially separates a portfolio's performance with respect to a benchmark into its upside U_{P_t} and its downside D_{P_t} and places a value on the upside and downside by assuming an arbitrage-free market. The value of the upside translates intuitively into the **call value** of the portfolio while the value of the downside translates intuitively into the **put value** of the portfolio. Note, significantly, that by placing a value on the upside and the downside, a forward-looking performance measure has been defined based strictly on the MtF values of existing securities.

These definitions of call value and put value as applied to a portfolio's future distribution enables the terminology and intuition of **option pricing** to be applied to portfolio valuation where the underlying process is driven by a potentially vast number of underlying risk factors.

By definition, the call value always equals the put value for a given portfolio in an arbitrage-free and frictionless market (assuming the current value of the benchmark is equal to that of the optimal portfolio). Therefore, a feasible frontier, as defined by the call value against the put value, is a 45° line (i.e. having a slope of one) and contains all fairly priced securities.

When investors have a personal risk adjustment factor λ equal to one, they possess the same level of risk aversion as the market in general. From their perspective, all securities are equally desirable. In other words, the call value always compensates appropriately for the put value.

However, when investors possess a λ that exceeds one, they are more risk averse than the market and, thus, prefer to hold positions that mimic the benchmark. If an

Market Risk	Variance	$\sigma_{it}^2 = \sum_{j=1}^{S_t} p_j \left(m_{ijt} - \left(\sum_{j'=1}^{S_t} p_{j'} m_{ij't} \right) \right)^2$
	Standard Value-at-Risk (confidence level α)	$VaR_{it}(\alpha) : \Pr\{(m_{i0t} - m_{ijt}) \geq VaR_{it}(\alpha)\} = 1-\alpha$
	Expected Shortfall (confidence level α)	$E[S_{it}](\alpha) = \frac{1}{1-\alpha} \sum_{j=1}^{S_t} p_j \left(m_{i0t} - m_{ijt} - VaR_{it}(\alpha)\right)^+$
	Expected Downside	$E[D_{it}] = \sum_{j=1}^{S_t} p_j (m_{ijt} - \tau_{jt})^-$
	Regret	$E[D_{it}]^* = \lambda \cdot \sum_{j=1}^{S_t} p_j (m_{ijt} - \tau_{jt})^-$
	Put Value (Value of Downside)	$Put_{it} = \left(\sum_{j=1}^{S_t} p_j (m_{ijt} - \tau_{jt})^- \right) m_{Ijt}$
	Expected Counterparty Credit Exposure	$E[TE_{it}] = \sum_{j=1}^{S_t} p_j TE_{ijt}$
Credit Risk	Expected Counterparty Credit Loss	$L_i = \sum_{j=1}^{S_t} p_j \cdot \sum_{t^*=t}^{T} AE_{ijt^*} \cdot p_i(t^*\|j) \cdot (1 - r_i(t^*\|j))$
	Expected Cross-Counterparty Credit Loss	$L_t = \sum_{j=1}^{S_t} p_j \cdot \sum_{t^*=t}^{T} \cdot \sum_{i=1}^{N} AE_{ijt^*} \cdot p_i(t^*\|j) \cdot (1 - r_i(t^*\|j))$
Liquidity Risk	Liquidation Period (t^*) Adjusted VaR (confidence level α)	$VaR^*_{it}(\alpha) : \Pr\{(m_{i0t^*} - m_{ijt^*}) \geq VaR^*_{it}(\alpha)\} = 1-\alpha$

Table 7.2: Selected risk measures by risk class

$p_i(t^*|j)$ and $r_i(t^*|j)$ represent the probability of default and recovery rate (given default) at time step $t = t^*$ under scenario j. The parameter α represents a one-sided confidence interval. τ_{jt} represents the MtF value of a benchmark portfolio at time t under scenario j. p_j represents the risk-neutral probability of state j occurring and m_{Ijt} is the zero coupon basis instrument associated with time step t. It is assumed that $j = 0$ represents the base scenario. The functions $(x)^+$ and $(x)^-$ equal max $(x,0)$ and (min $(x,0)$), respectively. The parameter λ represents a specified risk aversion coefficient.

investor's λ is constant for all levels of put value and if markets are infinitely liquid (no long or short constraints on the purchase of securities), the investor will hold only a portfolio that mimics the benchmark.

When investors possess a λ that is less than one, they are less risk averse than the market and prefer to hold positions that accept more put value in exchange for more call value. If λ is constant for all levels of put value and if markets are infinitely liquid, the solution is unbounded as these investors will continually leverage their portfolio to achieve more and more call value.

However, typically markets have finite liquidity and individual risk adjustment factors λ are likely to depend on the level of the put value. Thus, an optimal portfolio (with an optimal call value against put value) may indeed be determined under most conditions. If an individual (or an institution) is willing to accept a given amount of downside (in exchange for greater upside), an optimal portfolio may be determined that incorporates the impact of finite liquidity in the market. Individuals who are only willing to accept a small quantity of downside in search of greater upside may not be impacted greatly by finite market liquidity. In contrast, individuals who are willing to accept a greater amount of downside may encounter significant liquidity effects in their attempts to take on an appropriate portfolio position.

Reward	**Expected P&L**	$E_{t}[P\&L_{t}] = \sum_{j=1}^{S_t} p_j (m_{ijt} - m_{i0t})$
	Expected Return	$R_{it} = \sum_{j=1}^{S_t} \frac{p_j m_{ijt}}{m_{i0t}} - 1$
	Expected Upside	$E[U_{it}] = \sum_{j=1}^{S_t} p_j (m_{ijt} - \tau_{jt})^{+}$
	Call Value (Value of Upside)	$Call_{it} = \left(\sum_{j=1}^{S_t} p_j (m_{ijt} - \tau_{jt})^{+} \right) m_{tj}$
Performance	**Sharpe Ratio**	$S_{it} = \frac{R_{it}}{\sigma_{it}}$
	RAROC	$R^{*}_{it}(\alpha) = \frac{R_{it}}{VaR_{it}(\alpha)}$
	Risk-adjusted Upside (Expected Upside - Regret)	$U^{*}_{it} = E[U_{it}] - \lambda E[D_{it}]$
	Put/Call Efficient Frontier	$CP_{it} = Call_{it} - \hat{\lambda} Put_{it}$

Table 7.2: Selected reward and performance measures

(τ_{jt} represents the MtF value of a benchmark portfolio at time t under scenario j. It is assumed that $j = 0$ represents the base scenario. p_j represents the risk-neutral probability of state j occurring and m_{tjt} is the zero coupon basis instrument associated with time step t. The parameters λ and $\hat{\lambda}$ represent a specified risk aversion coefficient and risk adjustment factor, respectively. The function $(x)^{+}$ equals max $(x,0)$.)

TRADING OFF risk AND reward

A serious drawback in the mean-variance framework for trading off risk and reward is that it relies on spot and historical information only and ignores aging and settlement in the future. In this sense, it is not forward-looking. In contrast, a forward-looking framework for trading off risk and return would be based on quantities whose value is determined at some point in the future. It would also account for the effects of portfolio aging and, potentially, for the dynamics of the portfolio and would use scenarios that spanned the range of possible future events. Mark-to-Future (MtF) (Dembo 1998a, 1998b) is such a framework. In a Mark-to-Future world it can be shown that any portfolio may be disaggregated into its Put and Call value. A short position in the Put has the same payoff as the downside under all scenarios. The Call has the same payoff as the upside. By valuing the Put and Call, we can place a value today on the risk and reward of the portfolio for some given horizon. The difficulty is that the Put and Call are (complex) derivatives of many underlying factors, namely all the risk factors that affect the portfolio value. This paper shows how the Put and Call value may be determined in both complete and incomplete markets, and how for any portfolio risk and return in a forward looking world translates into the Put/Call trade-off. We call this trade-off the Put/Call Efficient Frontier. The Put/Call model, because it is based on MtF, will provide an appropriate trade-off between risk and reward for [illegible] market, credit, liquidity and other risks.

In a MtF framework (Dembo 1998a, 1998b), the upside can be modeled by a European Call option with maturity equal to the horizon and a strike given by the MtF value of investing an amount, equal to the mark-to-market value of the portfolio, in a benchmark instrument (or numeraire). A short position in a European Put option, with an identical strike and maturity, has the same payoff as the downside. The trade-off between risk and return for any portfolio is therefore captured by the trade-off between its Put and Call value. If Put/Call parity holds, the forward value of the deal is then (Call – Put). We define the risk-adjusted value to be (Call–λ^*Put) where λ is a parameter expressing an investor's aversion to risk*[1]. The linear function λ^*Put, which we call the investor's Regret, captures not only the downside but also the psychological factors that may influence an individual or an institution in their aversion to risk.

This is an inherently forward-looking view of risk-return since the value of the Put or the Call depends exclusively on future events; it is entirely oblivious of the past. Naturally, the past might influence one's choice of future scenarios, which will affect the Put or the Call value, but in all other respects the past is irrelevant.

How does one calculate the Put or Call value of an arbitrary portfolio? Moreover, in order for the Put/Call decomposition to be useful, we need to be able to value them in both complete and incomplete markets for portfolios of arbitrary composition. The model described in this chapter achieves these goals.

The Put/Call Efficient Frontier

The Put/Call Efficient Frontier traces the maximum upside for given levels of downside. Equivalently, it may be formulated as the minimum level of downside for fixed levels of upside. As we have shown, the upside in a MtF may be modeled and priced as a European Call option on the future value of the portfolio, with strike equal to the MtF value of a benchmark instrument and maturity equal to the horizon. The downside may be expressed and priced as a European Put option. If the maximum tolerable Put value (downside) is parameterized by k, then the efficient frontier, $e(k)$ is found by solving, for various levels of k:

$e(k)$ =	Maximize:	Call Value
	Subject to:	Put Value not exceeding a prespecified value, k.

Figure 8.1 The Put/Call Efficient Frontier

The Put/Call frontier is computed in two stages, each one involving an optimization model. First, to obtain the Put/Call values we compute appropriate state-price vectors (or risk-neutral probabilities) to discount the downside and upside components of the MtF vector. As we will see, this requires the solution of an optimization problem since we cannot assume that the market is complete. This formulation is related to the OPR model in Dembo (1995). The Put/Call frontier is then obtained by solving a second, very similar, optimization model.

In defining the Put/Call model we make the following assumptions:

- The market is illiquid; that is, there always exists some level of trading at which liquidity constraints become binding. At this level, the price of the security increases with the quantity traded.
- Market participants (i.e. an individual, a bank etc.) will operate under an additional constraint; namely, they will only trade within a pre-specified downside limit. In other words, participants select a maximum downside that they can tolerate (by default, their equity) and, as we will see, this will determine how they price risk.

Notation

For ease of presentation we only show the single horizon model and drop the subscript t that denotes the time step.

Let q be an N-dimensional vector of observed market values, one for each security $i = 1,\ldots,N$.

Let M be an S by N matrix whose columns are the MtF values of each security. S is the number of scenarios and p is a vector of prior (subjective) scenario probabilities.

If x is a portfolio of positions in these N securities, then $q^T x$ is the mark-to-market and Mx is the MtF of the portfolio. We assume that liquidity is modeled by allowing each security to be represented by a piecewise-linear function, with the price being a monotonic increasing function of volume. This then results in bounded positions, $\underline{x} \le x \le \bar{x}$ and instruments that are tranches of the actual traded securities. The bounds that this introduces into the model are not arbitrary trading amounts but forecasts of the price/volume relationship of each and every security. In principle, any amount of any security will be allowed into a solution if the model is willing to pay the price. The only true upper bound on the security might be the total number in circulation. We assume there is an absolute lower limit on the amount of shortselling for each security and that there is only one price at which shortselling occurs, regardless of the amount sold (i.e. one shortselling tranche). The lower bounds, $\underline{x}$, for all tranches of a traded instrument are all zero with the exception of the lowest tranche (short-selling) whose lower bound may be negative. The upper bounds $\bar{x}$ on tranches are all positive or zero. A security is then the sum of its tranches (see Figure 8.2)

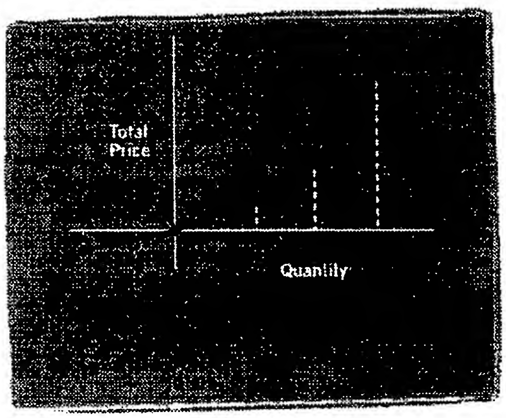

Figure 8.2 Modeling Liquidity: Total Price vs. Quantity is piecewise-linear

Let r be a column vector of dimension S whose elements are the value of a unit of the numeraire at the horizon for each and every scenario; then $Mx - rq^T x$ is a vector of future values of the portfolio minus the future value of the benchmark, that is, the net gain or loss (P&L) relative to holding an amount, equal to the mark-to-market value, in the numeraire.

Define $u(x) \equiv (Mx - rq^T x)^+ \ge 0$ to be the future P&L upside and $d(x) \equiv |(Mx - rq^T x)^-| \ge 0$ to be the absolute value of the future P&L downside. Then, by definition, $u(x) - d(x) \equiv Mx - rq^T x$.

Let $\rho \in R^S$ be a set of risk-neutral probabilities[3]. Then $\rho^T u(x)$ and $\rho^T d(x)$ are Call and Put values respectively, on the future P&L.

Formulating the Put/Call Efficient Frontier Model

In the formulation below, we allow u and d to be implicit functions of x. That is, they are treated as variables. The Put/Call Efficient Frontier is then generated by solving the following linear program (LP) for various levels of expected downside, k (see Dembo 1995).

$$e(k) \equiv \text{Maximize}_{(x,u,d)}: \ \rho^T u$$

$$\begin{array}{ll} \text{Subject to:} & \rho^T d \le k \\ & u - d - (M - rq^T)x = 0 \\ & \underline{x} \le x \le \bar{x} \\ & u \ge 0; \quad d \ge 0 \end{array}$$

The upper and lower bounds on the positions, x, arise solely out of liquidity considerations. The different prices for different tranches of securities made up by subsets of x are reflected in the mark-to-market, q, and MtF values, M.

One of the difficulties of this formulation is that a risk-neutral probability vector, ρ, is not known a priori. To obtain ρ we first solve a related problem, with the Call and Put being replaced by the expected upside, $p^T u(x)$, and expected downside, $p^T d(x)$, respectively (i.e. replace ρ with p). Once this has been solved ρ can be recovered and the true Put/Call Efficient Frontier generated.

$$e_p(k) \equiv \text{Maximize}_{(x,u,d)}: \ p^T u$$

$$\begin{array}{lll} \text{Subject to:} & p^T d \le k & (\mu) \\ & u - d - (M - rq^T)x = 0 & (\pi) \\ & \underline{x} \le x \le \bar{x} & (\underline{\omega});(\bar{\omega}) \\ & u \ge 0; \quad d \ge 0 & \end{array}$$

We refer to $e_p(k)$ as the **Primal** problem. The values in parentheses associated with each constraint are the dual variables.

The **Dual** of this problem, $e_d(k)$, is:

$$e_d(k) \equiv \text{Minimize}_{(\mu, \pi, \underline{\omega}, \bar{\omega})}: \ k\mu - \underline{x}^T \underline{\omega} + \bar{x}^T \bar{\omega}$$

$$\begin{array}{lll} \text{Subject to:} & (M - rq^T)^T \pi + \underline{\omega} - \bar{\omega} = 0 & (x) \\ & -\pi + p\mu \ge 0 & (d) \\ & \pi \ge p & (u) \\ & \mu \ge 0; \quad \underline{\omega} \ge 0; \quad \bar{\omega} \ge 0 & \end{array}$$

The values in brackets associated with these constraints are the primal variables.

This is a linear programming primal/dual pair with the following data requirements:

- $\mathbf{M}$ MtF of eligible securities
- p (subjective) probabilities associated with MtF scenarios
- q Mark-to-Market/Model (MtM) or current price
- r the growth rate of the numeraire over the horizon under each scenario
- $\underline{x}$; $\bar{x}$ lower and upper ranges for which MtM price is valid[4].

At the solution to this LP the primal and dual objective functions must be equal. This provides the equation describing the (expected) upside/downside efficient frontier. Equating the two objective functions gives:

$$e_p(k) = k\mu - \underline{x}^T \underline{\omega} + \bar{x}^T \bar{\omega}$$

If the upper and lower bounds on x are either not present or inactive at a solution, then $e_p(k) = k\mu$, which is the upside to downside trade-off when there is infinite liquidity. As these constraints kick-in, they reduce the maximum upside. As we allow for more and more downside (i.e. increase k) more and more of these will become active and their shadow prices $(\underline{\omega}, \bar{\omega})$ will change. Notice that $(-\underline{x}^T \underline{\omega} + \bar{x}^T\bar{\omega})$ is always non-negative, however, since $\underline{x}$ is zero or negative, $\bar{x}$ is zero or positive and $\underline{\omega}, \bar{\omega} \geq 0$. This implies that $e_p(0) \geq 0$. Thus the efficient frontier does not necessarily touch the origin.

The dual prices μ, $\underline{\omega}$ and $\bar{\omega}$ change only at certain values of k[5]. Thus, $e_p(k)$ is a piecewise-linear function with slope μ (i.e., $\mu = \partial\, e_p(k)/\partial k$). Moreover, it can be shown that μ is a monotonic decreasing function of k, and so the efficient frontier is concave (Figure 8.3). This is consistent with the fact that, as downside increases, the most attractive securities (or tranches) reach their trading limits and we are forced to enter into ever-inferior positions[6].

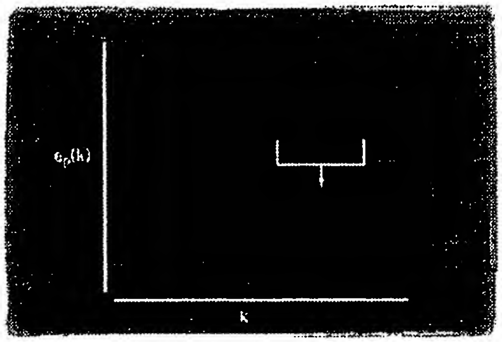

Figure 8.3 Upside/Downside Efficient Frontier (Solution to LP Model)

The complementarity conditions for this Primal/Dual linear programming pair are:

$$\begin{aligned}
\pi^T(u - d - (\mathbf{M} - rq^T)\, x) &= 0\\
(p^T d - k)\mu &= 0\\
\bar{\omega}^T(x - \bar{x}) &= 0\\
\underline{\omega}^T(\underline{x} - x) &= 0\\[1em]
x^T((\mathbf{M} - rq^T)^T \pi + \underline{\omega} - \bar{\omega}) &= 0\\
d^T(-\pi + p\mu) &= 0\\
u^T(\pi - p) &= 0
\end{aligned}$$

At optimality the primal variables x, u, d and the dual variables π, μ, $\underline{\omega}$, $\bar{\omega}$ must satisfy primal feasibility, dual feasibility and complementarity conditions.

Dual feasibility requires $\mu \geq 1$. From complementarity, this implies $p^T d = k$ in any optimal solution. Notice also that $\mu = 1$ implies $\pi = p$ for any dual-feasible solution. If $\pi = p$ is feasible in the dual then, as we will show below, p is a risk-neutral probability vector. If $\pi = p$ is not dual-feasible, this implies $\mu > 1$.

Dual feasibility also implies that

$$r_0 \mathbf{M}^T \rho = q + (r_0 / \Sigma_j\, \pi_j)\, (\bar{\omega} - \underline{\omega})$$

where $\rho = \pi / \Sigma_j\, \pi_j$ is a vector of risk-neutral probabilities and $r_0 = 1 / r^T \rho$ is a risk-neutral discount factor. The dual prices $\underline{\omega}$ and $\bar{\omega}$ give rise to liquidity discounts and premiums, respectively, that properly price securities in a risk-neutral framework. Thus, we can view the quantity

$$q^{adj} = q + (r_0 / \Sigma_j\, \pi_j)\, (\bar{\omega} - \underline{\omega})$$

as a set of liquidity-adjusted risk-neutral security prices that satisfy

$$r_0\, \mathbf{M}^T \rho = q^{adj}$$

We can interpret the term $(r_C / \Sigma_j\, \pi_j)(\bar{\omega} - \underline{\omega})$ as the amount by which today's prices need to be adjusted to reflect

market liquidity. The probability vector ρ may then be thought of as *liquidity-adjusted risk-neutral probabilities.* If $\hat{\omega}_j > 0$ then $\underline{\omega}_j = 0$ and there is a liquidity premium that we apply to the mark-to-market value q_j in order to compute appropriate risk-neutral discount factors. If, on the other hand, $\omega_j = 0$ *and* $\underline{\omega}_j > 0$ then the mark-to-market value is too high and we apply an effective discount to q_j. The larger these dual prices, the more the market wants to push through the barrier imposed by the liquidity bound and, therefore, the higher the premium/discount. If the liquidity model accurately reflects market price behavior then this information properly calibrates risk-neutral probabilities to obtain more realistic prices.

The first complementarity condition implies

$$r_0\, \rho^T u - r_0\, \rho^T d = r_0\, \rho^T (M - r q^T) x$$

which is a statement of *Put/Call Parity,* since $r_0\, \rho^T u$ is the value of a Call and $r_0\, \rho^T d$ is the value of a Put on the future P&L with respect to the numeraire. Thus, π, and hence ρ, which are obtained as by-products of solving the above linear program, are the missing link needed to properly price the Put and Call in a MtF framework.

The majority of the work required to set up and solve the Put/Call model is in computing the MtF matrix, **M**. Once this has been done, the model can be easily solved using standard LP software. Even problems with many thousands of instruments and scenarios can be solved in a few minutes on a PC.

Finally, once the risk-neutral probabilities, ρ, have been calculated we are in a position to compute the true Put/Call efficient frontier. We do this by solving the original model $\epsilon(k)$, which is effectively equivalent to solving $\epsilon_\rho(k)$ with ρ in place of the subjective probabilities p. All of our preceding observations regarding the upside/downside efficient frontier, namely its piecewise-linearity and concavity, extend naturally to the Put/Call Efficient Frontier. The key difference is that the Put/Call Efficient Frontier has been calibrated to reflect the implied risk preferences of investors.

Note that each segment of the upside/downside efficient frontier potentially gives rise to a different set of risk-neutral probabilities. Thus, investors with different tolerances for risk (as measured by the expected downside, k) may in fact obtain different risk-neutral probabilities, discount factors and Put/Call Efficient Frontiers. Figure 8.4, shows two investors, A and B, with respective risk tolerances k_A and k_B. Being on the second segment of the upside/downside efficient frontier, investor A uses the corresponding probabilities ρ_2 to construct a Put/Call Efficient Frontier. In contrast, investor B's greater risk tolerance places her on the third segment of the upside/downside efficient frontier, and her Put/Call Efficient Frontier is constructed using probabilities ρ_3.

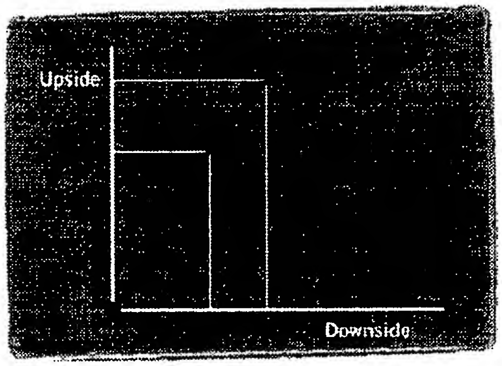

Figure 8.4 Different risk tolerances and risk-neutral probabilities

Recall that we use (Call – λ*Put) to measure risk-adjusted performance, where λ is a parameter expressing an investor's risk-aversion relative to the prevailing market preferences. An investor's λ determines which of the Put/Call efficient portfolios will be preferred. Specifically, an investor with risk-aversion λ_0 will select:

- an efficient portfolio with $\mu = \lambda_0$ or, if there is no such portfolio,
- an efficient portfolio with $\mu = \lambda_0$ and the largest call value

Linking Market, Credit and Other Risks

One of the most important aspects of the framework we have presented is that it is entirely generic. In developing the Put/Call Frontier we have made no assumptions about whether market, credit or any other risk is involved. If the scenarios contain events that link these types of risk and if the MtF is properly computed, then the Put/Call Frontier will be an efficient frontier that links all these risks and trades them off with their upside. Correlation between all these types of risks will be taken into account, since it is implicit in the choice of scenarios.

This further reinforces one of the primary principles behind MtF. By separating the simulation from the risk measure itself, the way in which we trade-off risk and return and link various types of risks may be developed in a generic manner.

[1] This section is based on Dembo [1999]. The MtF framework supports a wide range of models, including scenario optimization [Dembo 1991] and optimal portfolio replication [Dembo 1995].

[2] In the Put/Call framework, $\hat{\lambda}$ actually presents an investor's degree of risk aversion relative to that of the market in general (e.g., $\hat{\lambda} = 1$ for an investor that is as risk-averse as the market, while investors with $\hat{\lambda} < 1$ and $\hat{\lambda} > 1$, respectively, less and more risk-averse than the market).

[3] In fact, ρ is not necessarily a risk-neutral probability when the market is incomplete (see also: Dembo [1995], Dembo and Rosen [1999]). It is the appropriate discount factor we need for pricing a Put or Call in a MtF framework. In an incomplete market ρ may be viewed as a set of Regret-Neutral or Utility-Neutral probabilities, which may be used to price a Call or Put.

[4] We allow for liquidity effects by modeling the price of each security as a piecewise-linear function. This requires that each security be represented as a sum of bounded securities, with a particular order in which the securities are "filled". The MtF (M) and mark-to-market (q) values reflect the relative desirability of the securities and ensure that they are filled in the required sequence.

[5] Specifically, when there is a change in the optimal basis of the linear program.

[6] To allow proper interpretation of the results, we implicitly assume that $\mu > 1$ (i.e., market prices reflect the fact that we require more than one unit of upside to take on one unit of downside). This assumption can be relaxed if the (nonlinear) constraint $u^T d = 0$ is added to problem $e_p(k)$.

References

Artzner, P., F. Delbaen, J.M. Eber and D. Heath, 1998, "Coherent measures of risk," Working Paper, Institut de Recherche Mathématique Avancée, Université Louis Pasteur et C.N.R.S.

Bank for International Settlements, Amendment to the Capital Accord to Incorporate Market Risks, Basle Committee on Banking Supervision, January 1996.

Boyle, P., 1977, "New life forms on the option landscape," The Journal of Financial Engineering, 2(3): 217-52.

Cartolano, P. and S. Verma, 1999, "Using scenario banding to stress test counterparty credit exposures," forthcoming in Algo Research Quarterly, 2(4).

Chishti, A., 1999, "Simulation of fixed-income portfolios using grids," Algo Research Quarterly, 2(2): 41-50.

Credite Suisse Financial Products, 1997, CreditRisk+: A Credit Risk Management Framework, New York, NY.

Dembo, R.S., 1991, "Scenario optimization," Annals of Operations Research, 8: 267-284.

Dembo, R.S. and A. Freeman, 1998, Seeing Tomorrow: Rewriting the Rules of Risk, New York: John Wiley & Sons, Inc.

Dembo, R.S., 1999, "Optimal portfolio replication," Research Paper Series 95-01, Algorithmics, Inc.

Dembo, R.S., 1991, "Scenario Optimization," *Annals of Operations Research*, 30:63-80.

Dembo, R.S., 1995, "Optimal Portfolio Replication," Research Paper, Algorithmics Inc.

Dembo, R.S., 1998a, "Mark-to-Future: A Consistent Firm-Wide Paradigm for Measuring Risk and Return," in *Risk Management and Analysis, Vol 1: Measuring and Modeling Financial Risk*, Carol Alexander (ed.). New York: John Wiley & Sons Ltd.

Dembo, R.S., 1998b, "Mark-to-Future," Morgan Stanley Dean Witter Global Equity and Derivative Markets, March 11, 1998: 6-16.

Dembo, R.S., 1999, "The Put/Call Efficient Frontier," Working Paper, Algorithmics Inc.

Dembo, R.S. and D. Rosen, 1999, "The practice of portfolio replication: A practical overview of forward and inverse problems," *Annals of Operations Research*, 85:267-284.

Fuerbringer, J. 1999, "When being wrong won't hurt," The New York Times, October 15, 1999: C1

Garman, M.B. and S.W Kohlhagen, 1983, "Foreign currency options values," Journal of International Money and Finance, 2(Dec.): 231-37.

Iscoe, I., A. Kreinin and D. Rosen, 1999, "An integrated market and credit risk portfolio model," Algo Research Quarterly, 2(3): 21-37.

Jamshidian, F. and Y. Zhu, 1997, "Scenario simulation: theory and methodology," Finance and Stochastics, 1 (1): 43-67.

Kreinin I., L. Merkoulovitch, D. Rosen and M. Zerbs, 1998a, "Measuring portfolio risk using Quasi Monte Carlo methods," Algo Research Quarterly, 1(1): 17-25.

Kreinin I., L. Merkoulovitch, D. Rosen and M. Zerbs, 1998b, "Principal component analysis in Quasi Monte Carlo simulation," Algo Research Quarterly, 1(2): 21-29.

Longerstaey, J. and P. Zangari, 1996, RiskMetrics-Technical Document, 4th ed., New York: Morgan Guaranty Trust Co. Available at http://www.riskmetrics.com/rm/index.cgi

Locke, J., 1999, "A fear of floating," Latin Risk, (Risk Special Report), July: 4-5.

Markowitz, H.M., 1952, "Portfolio selection," The Journal of Finance, 7(1): 77-91.

Yung, E., 1999a, "Making a scene...," Algo Research Quarterly, 2(1): 5-7.

Yung, E., 1999b, "Making a scene...," Algo Research Quarterly, 2(2): 5-9.

Yung, E., 1999c, "Making a scene...," Algo Research Quarterly, 2(3): 5-8.

Part 3
Technology Considerations

A case study based on a hypothetical global bank is used to illustrate the need for a framework such as MtF. The many benefits of the distributed MtF process are compared to the limitations of traditional, centralized risk management systems. A case study based on an actual implementation of MtF at HypoVereinsbank is used to discuss practical considerations.

3

Enterpri

Mark-to-Future

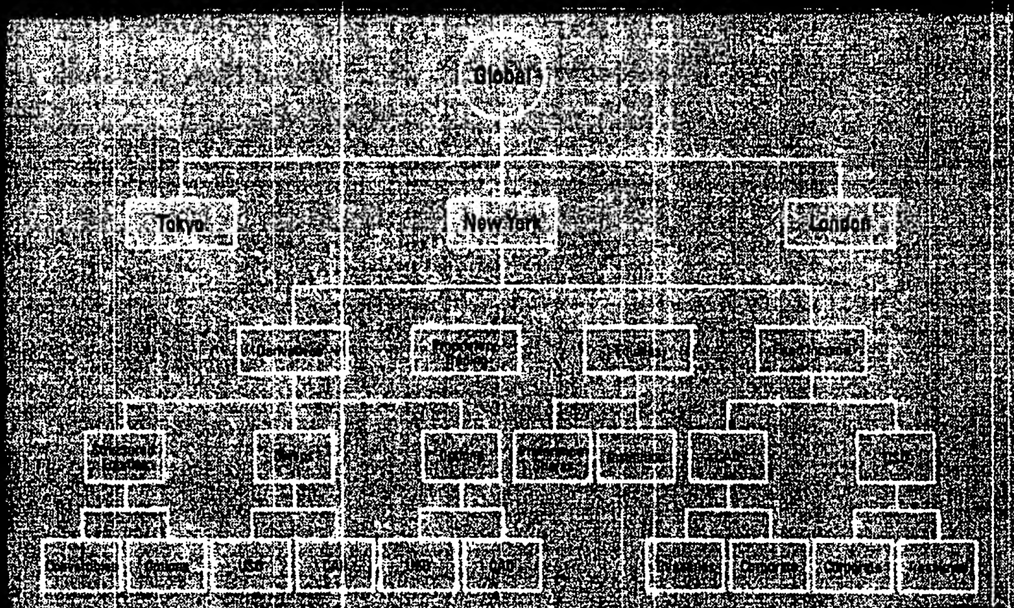

Figure 1.1: Market risk limits structure

This chapter describes a hypothetical bank and outlines the bank's risk reporting requirements with respect to its market risk and credit risk limit structures. The required risk measures are enumerated and the methods of calculating them described. An approach to the bank's reporting requirements based on a traditional, centralized architecture is considered. A review of the shortcomings of a traditional approach motivates the discussion of approaches based on the MtF framework. Centralized and distributed processes based on the MtF framework and the benefits realized by adopting this framework are presented.

se-Wide

Implementation

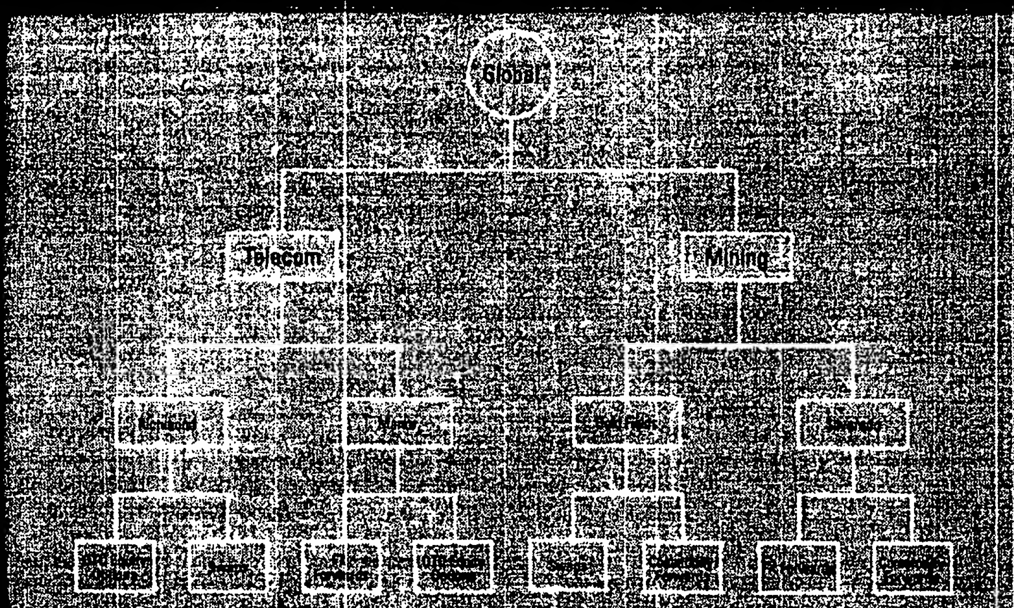

Figure 1.2: Credit risk limits structure

The objective of this chapter is to use the risk reporting requirements of a hypothetical financial institution, the Spadina Bank, to demonstrate the benefits of implementing enterprise risk management within the MtF framework. It does not purport to make recommendations on risk management best-practices.

The Spadina Bank is a New York-based investment bank with major trading operations in Tokyo, London and New York. It is a market maker in a variety of interest rate, equity and foreign exchange products—including both exchange-traded and over-the-counter (OTC) derivatives. It holds approximately 200,000 positions worldwide 50,000 each in Tokyo and London and 100,000 in New York. Of the 200,000 positions, 120,000 are in over-the-counter instruments; the remaining 80,000 are in 5,000 exchange-traded instruments.

The market risk limits structure of the bank presented in Figure 1.1 follows the organizational hierarchy of the bank, from the enterprise level through the branches and desks to the products at each desk.

The credit risk limits structure is organized by industry sector, counterparty and product type. The credit risk limits structure for two sectors, telecommunications and mining, is presented in Figure 1.2. The Spadina Bank deals with two counterparties in telecommunications, Richmond and Manor, in OTC equity options, swaps and FX forwards.

The primary objectives of the Spadina Bank's enterprise-wide risk management system are to

- deliver validated local risk reports and draft global risk reports by 09:00 local time
- satisfy regulatory requirements for risk reporting
- effectively allocate credit by calculating counterparty exposures that account for netting and portfolio diversification
- provide decision support to traders on a timely basis in order to maximize returns for a given level of risk capital.

In accordance with recent regulatory best-practice guidelines, the bank requires an integrated framework for market, credit and liquidity risk management. The local and global market risk reports must satisfy regulatory reporting requirements for institutions using internal models.

Risk Reporting Requirements

The bank's objectives for market, liquidity and credit risk management are outlined below.

Market Risk Measures

The Spadina Bank's primary market risk measures include a global non-parametric VaR at the 99% confidence level for a 10-day holding period, scenario-based sensitivities at the desk level and equity-specific risk. The bank relies on comprehensive stress scenarios for market and liquidity risk to assess losses under extreme and relevant market events.

Liquidity-Adjusted Global VaR. Global VaR is calculated using 1,000 multi-step Monte Carlo scenarios over a 10-day time horizon, using time steps of three different lengths: two days, five days and 10 days. There are four scenario sets: a set of interest rate scenarios; FX scenarios; equity risk scenarios; and a fourth set comprising a full set of joint interest rate, equity, and FX scenarios.

The bank classifies all positions into one of three categories—highly liquid, liquid and illiquid. This classification is reviewed on a quarterly basis. The highly liquid positions are settled into cash in two days, the liquid positions in five days and the illiquid positions are settled at the 10-day risk horizon. A dynamic portfolio strategy is required to capture the liquidity premium over the 10-day horizon.

Scenario-Based Sensitivity Analyses. The Spadina Bank uses 118 product-specific scenario-based sensitivity analyses to isolate the impact of various risk factor shocks on the value of the portfolio. These scenarios include shocks to node points on curves, FX rates and volatilities. The bank has chosen an **instantaneous price shock**: the portfolio positions remain static while the portfolio is revalued under each scenario (BIS 1996).

Stress Testing. The bank uses 50 stress test scenarios based on historical data to account for atypical events. These data are selected from periods of high market volatility and from those periods during which the bank experienced significant losses. User-defined scenarios may also be defined for local measures.

Stress tests include standard historical stress scenarios such as the 1987 crash, or the 1998 flight-to-quality period and proxy scenarios for emerging market positions, as described in Part 2, Chapter 3. The bank is concerned about *market-neutral* arbitrage positions that show little exposure under standard stress tests but might exhibit high exposure under less obvious stress conditions. Hence, a small representative benchmark portfolio is evaluated periodically over a large historical time series to identify additional context-dependent stress scenarios. To assess liquidity risk, stress scenarios that capture market events over three time steps are used.

Equity-Specific Risk. The Spadina Bank has chosen to use an internal model to measure equity-specific risk. The model is based on a simulation that uses two years of historical data.

Credit Risk Measures

The Spadina Bank must measure netted, global counterparty credit exposure. The bank uses 1,000 multi-step Monte Carlo scenarios, with 32 time steps, over 20 years. The bank conservatively assumes that the credit risk of OTC positions cannot be hedged over the life of the contract; hence, the horizon for the analysis is the maturity of the longest dated contract, in this case 20 years. The time steps are increasingly large: the bank uses daily steps for the first week, weekly steps for the first month, quarterly steps for the first year and yearly steps for the remaining years. The portfolio is aged through time accounting for both cash and physical settlement. MtF values are post-processed to compute counterparty credit exposures, both on a gross and netted basis using a 99% confidence level.

Summary of Reporting Requirements

The number of revaluations required to produce the risk measures described above for the global portfolio are summarized in Table 1.1. If no basis instruments are defined approximately 6.4 billion instrument valuations are required to produce the global risk report. For example, to calculate global VaR, 200,000 positions (120,000 OTC and 80,000 exchange-traded positions) are valued across 1,000 Monte Carlo scenarios and over three time steps. To calculate credit exposure, 120,000 OTC

Total VaR	200	1,000	3	600,000
Equity VaR	200	1,000	3	600,000
IR VaR	200	1,000	3	600,000
FX VaR	200	1,000	3	600,000
Scenario-based sensitivities	200	118	1	23,600
Stress testing	200	50	3	30,000
Equity-specific risk	200	500	1	100,000
Credit exposure	120	1,000	32	3,840,000
Total				[illegible]

Table 1.1: Summary of computations required for risk reporting measures with no basis mapping

instruments must be valued over 1,000 Monte Carlo scenarios and 32 time steps.

An efficient valuation engine can revalue between 500,000 and 2 million instruments per minute on a single 300 Mhz processor, depending on the complexity of the instruments. To be conservative, all time estimates are based on a revaluation rate of 500,000 instruments per minute. Approximately 210 hours (6.4 billion revaluations at 500,000 revaluations per minute) are required to compute the risk measures summarized in Table 1.1.

Clearly, to produce these risk measures within an overnight window the calculations must be performed in parallel. Parallel computations can be executed in a centralized or distributed process, within or outside of the MtF framework. Both centralized and distributed approaches are considered in the following sections. In the next section, a traditional, centralized approach that is not based on the MtF framework is presented. As the discussion demonstrates, the length of the overnight processing window is such that trade-offs must be made between the degree of parallelization and the degree to which the bank achieves its reporting objectives (Penny 1999).

A Centralized Approach Outside the MtF Framework

The non-additivity of most key risk measures contributes to many of the difficulties inherent in traditional enterprise risk management. For example, the credit exposure at a certain confidence level of a portfolio containing two derivatives is not the sum of the credit exposures of those two derivatives computed independently. In fact, no function of two independently computed credit exposures yields the portfolio credit exposure. This is equally true of credit losses, VaR and any type of risk-adjusted return.

The traditional approach to producing the global report is to consolidate end-of-day positions from each branch in the New York head office, to compute the required measures at the various portfolio levels in the hierarchy and then to process and produce the global reports centrally. Using this approach, the system must be centralized. For a global financial institution such as the Spadina Bank, this means collecting into a central data model all the terms and conditions and position holdings of all securities held in the enterprise, applying the appropriate market data and configuring appropriate pricing functions to perform valuations. It also results in a duplication of effort. The expertise required to deal with the various financial products distributed among the business units in Tokyo, London and New York must be duplicated at the head office.

The daily processes at the branches and head office of a traditional, centralized approach are illustrated in Figure 1.3. The three timelines display the local time in Tokyo, London and New York on a 24 hour clock. Alongside each timeline, horizontal bars illustrate the operations over time in each of the branches and in the New York head office. For the purpose of this discussion, the New York branch daily process is considered to be independent of the New York head office.

The daily processing in the Tokyo, London and New York branches is similar, however, the length of processing time required depends on the size and composition of the positions and the computing power at each branch. The processes at the Tokyo branch serve as an example for discussion of the components of the daily process at each site. Unless otherwise specified, all times are local times.

Prepare Market Data: The Tokyo market closes at 15:00 on day T. The Spadina Bank has made a policy decision to use Tokyo market data for Tokyo risk reports and for global risk reporting of the Tokyo positions. Accordingly, the closing market prices are used to bootstrap all required interest rate curves, credit spread curves and volatility surfaces. By 15:30 the bootstrapped, scrubbed market data for Tokyo is available.

Process Back Office and Transmit Data: By 22:00 on day T, the Tokyo end-of-day position data is assembled. At this stage, the position data is unvalidated—it has not been revised to reflect the errors and omissions that might be found at 07:00 on day $T+1$. Both the market data and the draft position data are transmitted to the head office, arriving at 08:00 EST (Eastern Standard Time) on day T in New York. The London market data and draft position data arrive in New York at 17:00 EST on day T and the New York draft data is assembled by 22:00 EST on day T.

Prepare Stress Scenarios: At 19:00 EST, before overnight processing begins, stress test and sensitivity scenarios for the next day are prepared centrally. A special risk committee, with input from senior management, prepares a consolidated and consistent scenario set. Most of the scenarios do not change day-to-day.

Prepare Draft Global and Local Reports: Outside of the MtF framework, the global position for each trading day, T, must be assembled centrally before processing begins. At 22:00 EST on day T the draft global position is assembled at the New York head office. There is a nine hour window for overnight processing. Given that 210 hours of computing time are required, 27 processors perform the risk calculations in parallel to produce the reports for the market risk and credit risk limits structures illustrated in Figures 1.1 and 1.2.

The draft local and global reports are transmitted to the branch offices at 22:00 on day $T+1$ in Tokyo, at 13:00 on day $T+1$ in London and at 08:00 on day $T+1$ in New York.

Errors and Omissions: At 07:00, the Tokyo risk manager arrives at work. Her first task is to review last night's risk reports and to examine e-mails from operations providing notification of data errors. On an average day, three to five of the 50,000 positions have either been omitted or entered incorrectly. These errors often significantly alter the risk profile at the desk level and occasionally alter the risk profile at the branch level. Errors can include missing instruments, incorrect positions, incorrect or missing static data, incorrect or missing market data and terms and conditions data.

The errors and omissions reports from the branches are transmitted to the New York head office as they are

Figure 1.3: Daily process in the centralized approach outside of the MtF framework

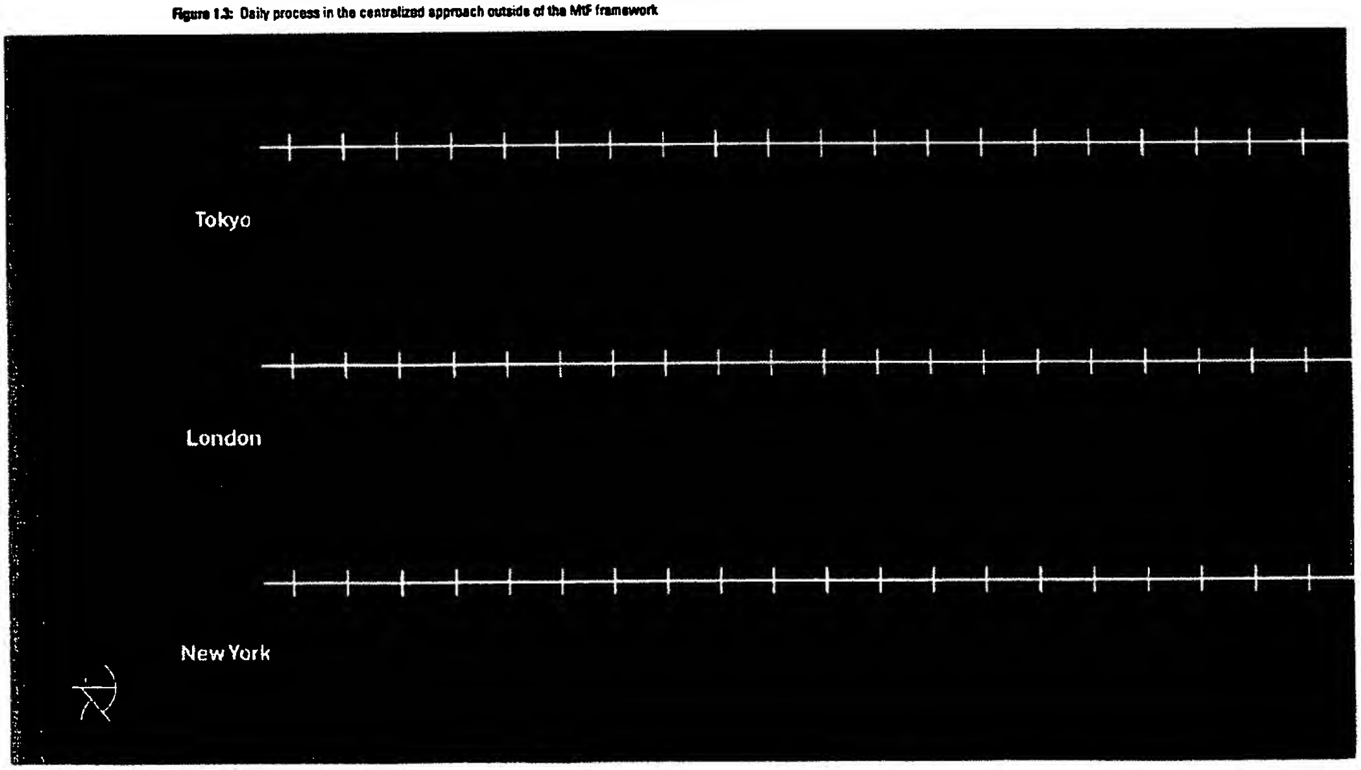

completed. The centralized data is valid after corrections are made based on the last errors and omissions report received from the New York branch at 08:00 EST on day $T+1$.

Prepare Final Global Report: At 08:00 EST on day $T+1$, the validated global position is assembled at the New York head office. A second nine hour processing window is required to produce the validated global risk report.

The final local and global reports are transmitted to the branch offices at 07:00 on day $T+2$ in Tokyo, at 22:00 on day $T+1$ in London and at 17:00 on day $T+1$ in New York. Each branch office has a validated local and global report for day T at the start of trading on day $T+2$.

Though workable, this is a complex implementation owing to the difficulties inherent in massive centralization. Most importantly, the minimum total processing time cannot be less than the time required to assemble all portfolio information in the central location. For a medium-sized bank such as this, collecting portfolio information centrally consumes the overnight processing window.

Thus, only New York receives draft global and local reports by 09:00 on day $T+1$. Validated global and local reports are not available until the start of trading on day $T+2$. Some of these difficulties can be mitigated by processing local positions to produce local reports. By duplicating the computing power of the head office across the branch offices and running local overnight processes, each branch can produce a draft, local report by the start of the trading day on $T+1$ and a validated, local report by mid-day on day $T+1$.

This centralized approach also has the disadvantage that reports must be pre-configured. Thus, a re-configuration of the set up and a new overnight computation or a doubling of the computing power is required to perform drill-downs or re-aggregations.

Centralizing all information is especially problematic when we consider the key incremental measures required to run a financial institution. Most prudent managers insist that decision makers understand the impact of their local decisions on the key global measures used to run the organization. In the case of revenue, this is straightforward: an additional \$1 earned locally is an additional \$1 earned globally. However, an additional \$1 of credit exposure incurred in a trading portfolio does not necessarily mean an additional \$1 of credit exposure at the enterprise level, because of netting, offsetting correlation effects, and the like.

Thus, a centralized approach is impractical as a method for feeding incremental, pre-deal analysis numbers back to the decision makers. Imagine a swaps dealer in Tokyo wishing to understand the impact of a proposed new deal on enterprise VaR and global credit exposure with a particular counterparty. The local application system must capture and transmit the details

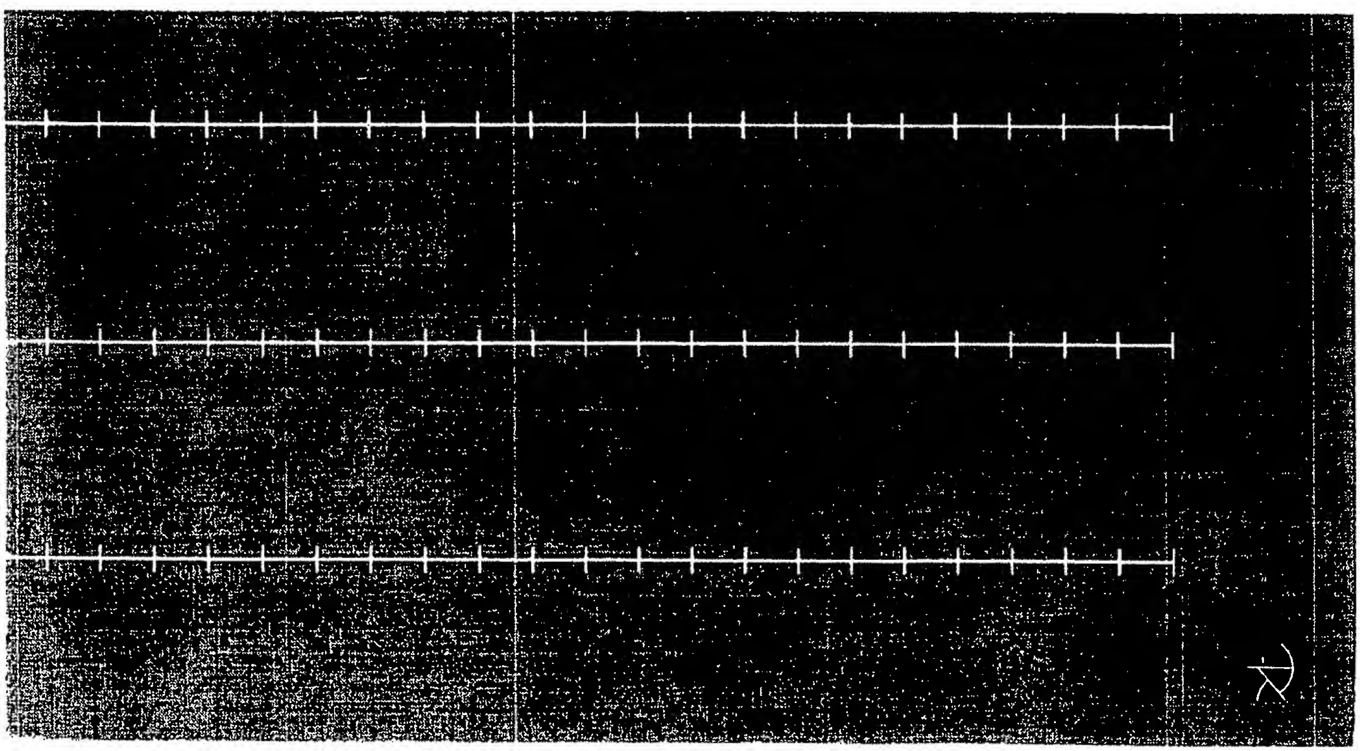

Zero coupon bond	50	50	500
Spot currency	49	1	49
Equities and indices	351	[illegible]	351
Commodities	20	5	100
Total	[illegible]	[illegible]	2000

Table 1.1: Details of mapping standard OTC positions to basis instruments

of a proposed deal to the central system at the head office in New York. A computation involving the tentative deal and all other open positions in the bank is then performed. Finally, the incremental numbers are returned to the Tokyo desk. Given the time zone differences, the most current global information is based on yesterday's draft data or validated data from two days earlier.

In summary, a centralized approach implemented outside of the MtF framework is workable but does not allow the Spadina Bank to achieve its reporting goals. In the next sections, we consider two approaches that exploit the additivity property of the MtF framework. The first is a centralized approach implemented in the MtF framework which permits the daily processes to be pipelined. The second is a distributed approach which creates further opportunities to improve performance. Daily processes implemented in the MtF framework allow the bank to achieve its risk management objectives.

A Centralized Approach Within the MtF Framework

A very significant reduction in processing time and storage of results is achieved by mapping financial products, particularly linear instruments such as bonds or swaps and plain-vanilla options, onto basis instruments as described in Part 2, Chapter 2. The 80,000 exchange-traded positions are mapped onto 5,000 unique exchange-traded products that serve as basis instruments. The 102,000 standardized OTC positions are mapped onto an additional 2,000 basis instruments; details of the mapping of these OTC positions is summarized in Table 1.1. The remaining 18,000 non-standardized OTC derivatives are mapped one-to-one onto basis instruments.

The mapping significantly reduces the number of instrument revaluations required for exchange-traded

Figure 1.4: Daily process in the centralized approach in the MtF framework

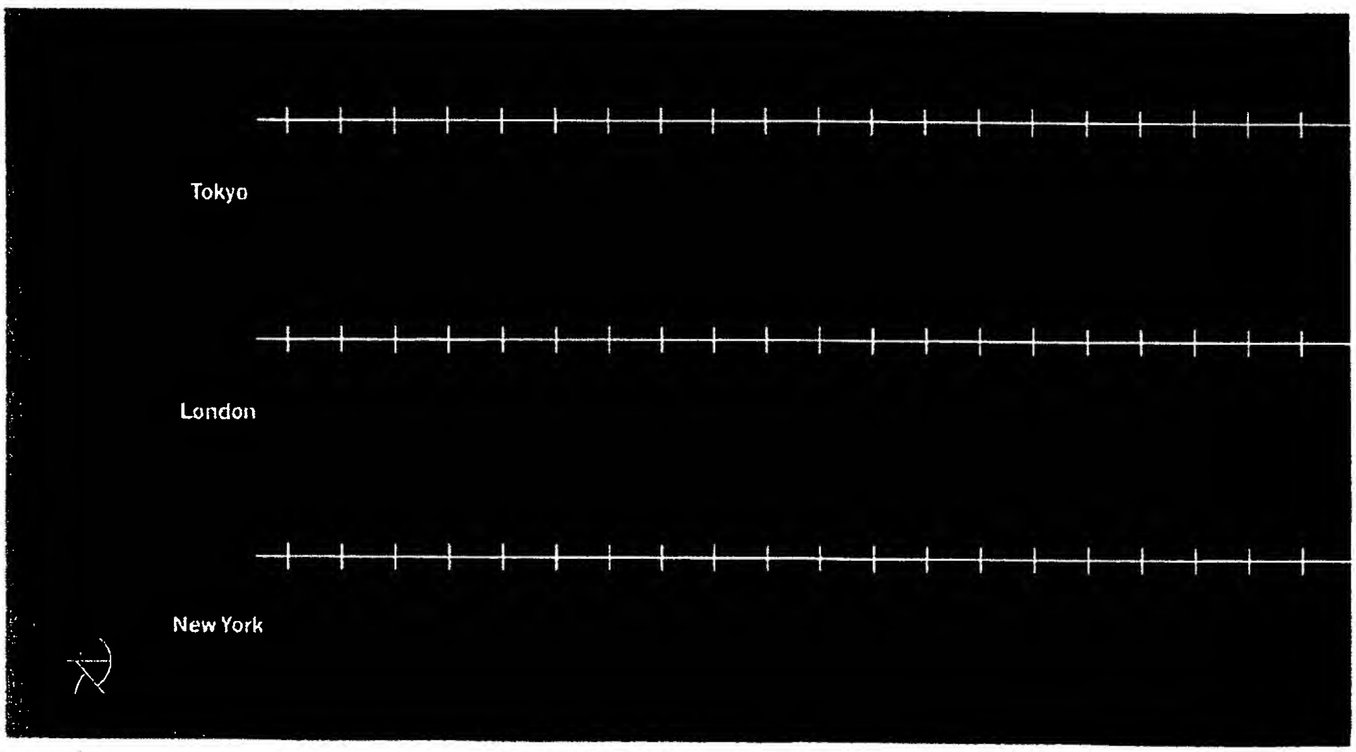

Global portfolio	200	100	25	87
Exchange-traded	80	40	5	94
OTC	120	60	20	83
Standardized OTC	102	51	2	98
Non-standardized OTC	18	9	18	0

Table 1.2: Percentage reduction realized by mapping onto basis instruments

Total VaR	25	1,000	3	75,000
Equity VaR	25	1,000	3	75,000
IR VaR	25	1,000	3	75,000
FX VaR	25	1,000	3	75,000
Scenario-based sensitivities	25	118	1	2,950
Stress testing	25	50	3	37,500
Equity-specific risk	25	500	1	12,500
Credit exposure	20	1,000	32	640,000
Total	[illegible]	[illegible]	[illegible]	959,200

Table 1.3: Summary of reduced computations required for risk reporting measures with basis mapping

and standard OTC positions (94% and 98%, respectively). No reductions are achieved for the non-standard OTC positions. The net reduction for the global portfolio is 87%. The percentage reduction by position type is summarized in Table 1.2.

The number of revaluations required to produce the risk measures described above for the global portfolio are summarized in Table 1.3. The total number of revaluations is reduced from approximately 6.4 billion to 1 billion and the processing time required is correspondingly reduced to approximately 32 hours.

The daily processes at the branch and head offices in

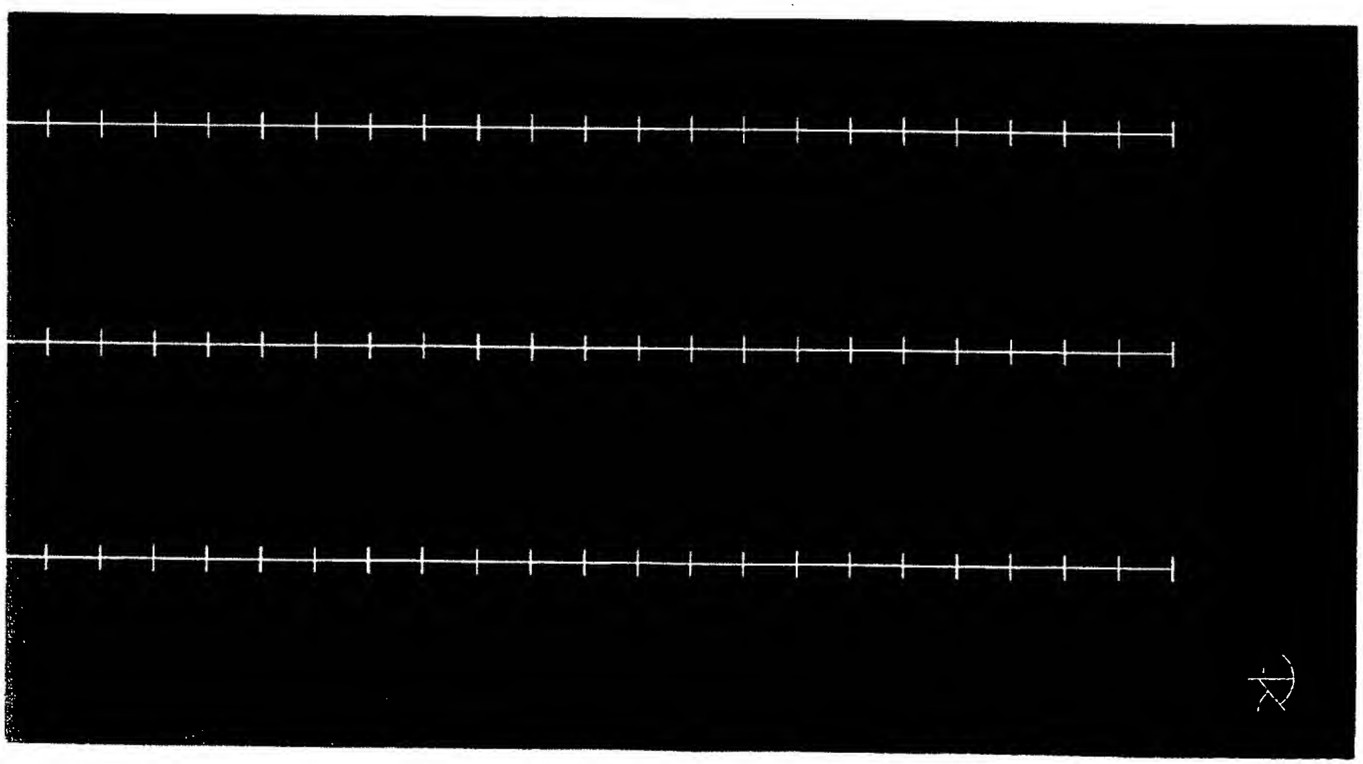

a centralized approach in the MtF framework are illustrated in Figure 1.4. Because processing is centralized, back office position data and market data are transmitted to the head office where it is used to generate the Global Portfolio MtF Cube according to the two hierarchies illustrated in Figures 1.1 and 1.2. This logical Cube may be represented by one or several physical Cubes.

Prepare Market Data, Process, Validate and Transmit Back Office Data: The daily process of preparing market data and back office processing and validation of position data is unchanged. However, because MtF values are additive (if the scenarios and time steps are consistent for each instrument, position and portfolio) processing of the branch data to generate Portfolio MtF Cubes can begin as soon as the position and market data are assembled. This is known as pipelining. Pipelining has two significant impacts: it reduces the need for massive parellelization and it extends the length of the window for overnight processing.

Generate Tokyo Basis MtF Cube: Because the window is wider, processing of the Tokyo position data can begin in New York at 18:00 EST on day *T*, after the correction of errors and omissions in Tokyo.

With 50,000 active positions, Tokyo manages 25% of the bank's global portfolio; accordingly, processing the Tokyo positions consumes 25% of the required 32 hours of processing time, or eight hours. To reduce processing time further, several MtF engines are run in parallel, each processing a sub-set of the basis instruments across the relevant scenarios and time steps.

Generate Portfolio MtF Cube: The Basis MtF Cubes are mapped into Product MtF Cubes and finally combined into Portfolio MtF Cubes across both the market and credit risk hierarchies to support timely *what-if* analysis during the trading day. The validated Tokyo Portfolio MtF Cube is transmitted to the Tokyo office at 11:00 local time on day *T+1*.

Generate Tokyo Risk Reports: Using the risk reporting application, the Portfolio MtF Cube is post-processed to calculate the risk measures of interest, summarized in Table 1.3. As these calculations are trivial—retrieving data, vector additions, sorting and determining confidence levels—the reports are generated in only 60 minutes. Local risk reports according to both the market and credit risk limits structures are available at 12:00 on day *T+1*.

Validated MtF Cubes created in the overnight process described above include validated MtF values for each instrument in the Tokyo portfolio, across all scenarios and time steps. The MtF system also stores the closing market data, from which the mark-to-market values for each instrument were calculated.

Processing of the London position can begin at 03:00

Figure 1.5: Daily process in a distributed approach in a MtF framework

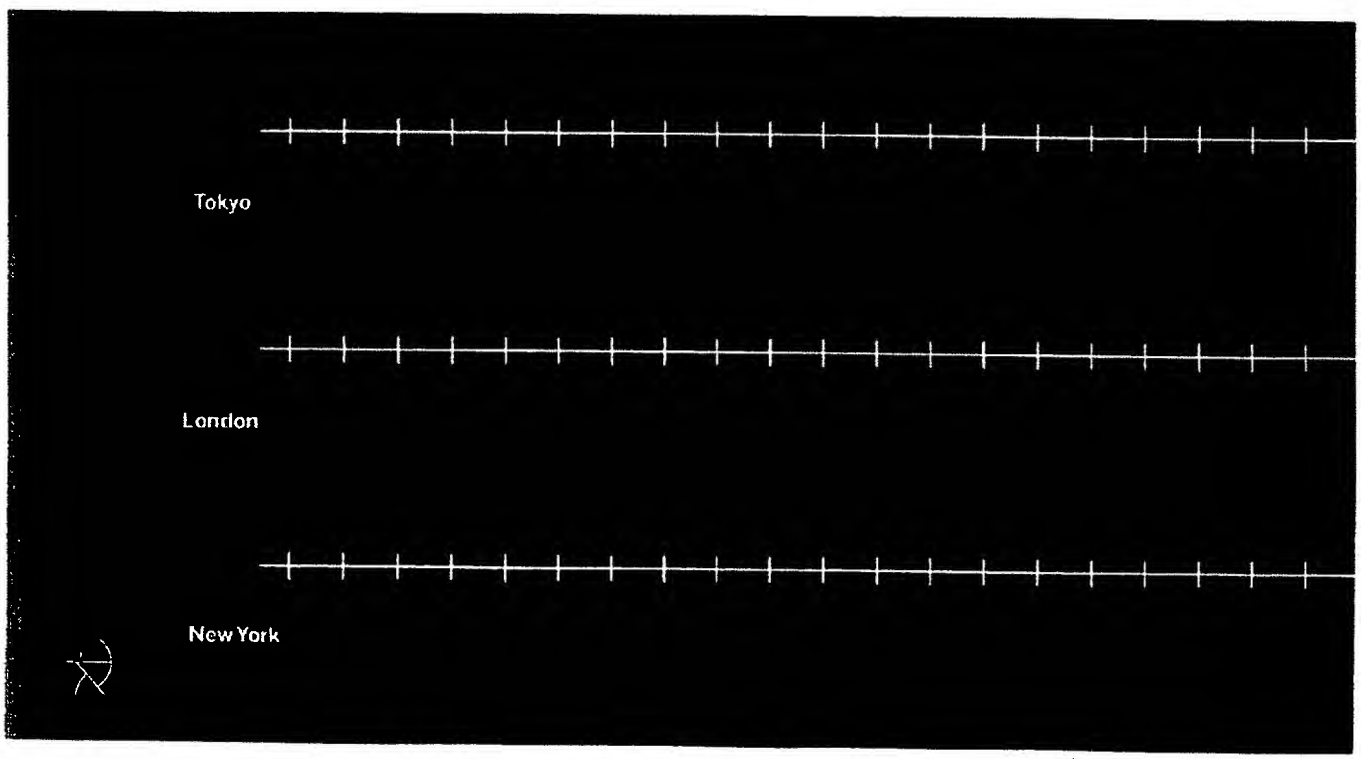

Prepare Tokyo market data	15:00	06:00	15:30	06:30
Process Tokyo back office	16:00	07:00	22:00	13:00
Generate Tokyo MtF Cube	16:00	07:00	23:00	14:00
Post-process and generate draft reports	23:00	14:00	00:00	15:00
Errors and omissions	07:00	22:00	08:00	23:00
Generate validated local risk reports	08:00	23:00	09:00	00:00

Table 1.4: Tokyo branch daily process at local time and Greenwich Mean Time

EST on day $T+1$ and processing of the New York position at 08:00 EST on day $T+1$. London has local reports at 12:00 on day $T+1$, one hour after the London MtF Cube is received. The New York local report is available one hour after the New York MtF Cube is available, at 14:00 on day $T+1$.

Produce and Distribute Global Risk Reports: The Spadina Bank global reports are available one hour after the New York MtF Cube is available, at 14:00 EST on day $T+1$.

The objectives of the Spadina Bank are more closely met by a centralized MtF approach: the full set of reports can be produced using fewer computing servers, but not by 09:00 local time. Given these shortcomings, the Spadina Bank has implemented a distributed process, which allows it to process MtF values in both a timely and an accurate manner.

A Distributed Approach Within the MtF Framework

Both centralized and distributed approaches can be implemented in the MtF framework, however a distinct advantage of MtF is that computations can be decentralized. In this section, a distributed approach to enterprise risk measurement is presented.

The daily processes at the branch and head offices of a distributed approach in the MtF framework are illustrated in Figure 1.5. Using this approach, local processing begins in a branch office when the market data is available, *before* the positions at the end of the trading day, T, are validated. The additivity of MtF values is exploited to update the position data and to correct errors and

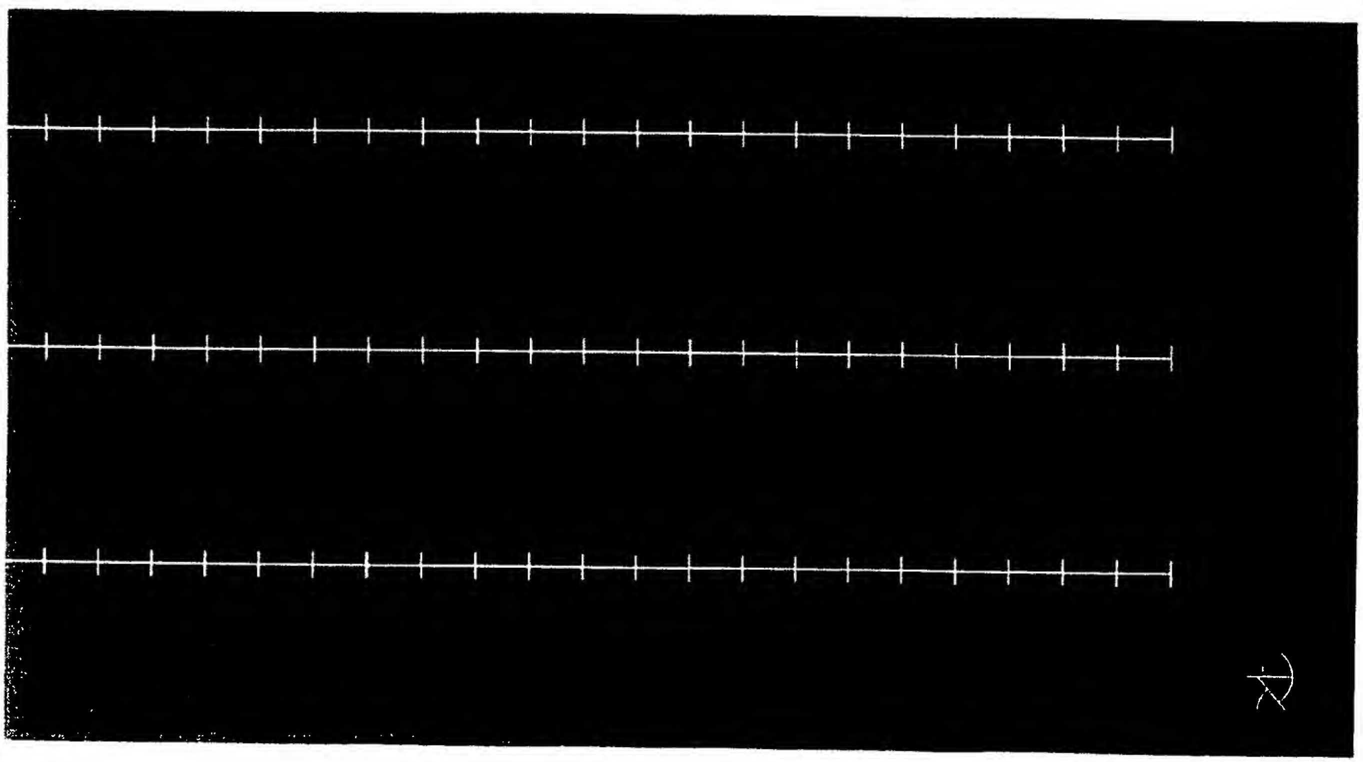

Prepare London market data	16:30	16:30	17:00	17:00
Process London back office	17:00	17:00	22:00	22:00
Generate London MtF Cube	17:00	17:00	23:00	23:00
Post-process and generate draft reports	23:00	23:00	00:00	00:00
Errors and omissions	07:00	07:00	08:00	08:00
Generate validated local risk reports	08:00	08:00	09:00	09:00

Table 1.5: London branch daily process at local time and Greenwich Mean Time

Prepare New York market data	17:00	22:00	17:30	22:30
Process New York back office	18:00	23:00	22:00	03:00
Generate New York MtF Cube	18:00	23:00	23:00	04:00
Post-process and generate draft reports	23:00	04:00	00:00	05:00
Errors and omissions	07:00	12:00	08:00	13:00
Generate validated local risk reports	08:00	13:00	09:00	14:00

Table 1.6: New York branch daily process at local time and Greenwich Mean Time

omissions. The processing window is widened sufficiently that all of the required risk reporting requirements summarized in Table 1.3 are performed and risk reports across all reporting hierarchies are delivered by 09:00 local time on day *T+1*.

Daily Processing at a Branch

Again, the daily processing in the Tokyo, London and New York branches is similar; the processes at the Tokyo branch illustrate the components of the daily process at each site. The majority of the processing occurs at the branches in the centralized approach. Therefore, the branch daily process is specified more explicitly in the discussion of the distributed approach.

Tokyo Branch Processing. The Tokyo branch daily process is summarized in Table 1.4.

Prepare Market Data and Process Back Office Data: The daily processes of preparing market data and of processing back office position data are unchanged.

Generate Tokyo Draft Basis MtF Cube: MtF processing begins at 16:00 on the draft portfolio, positions in OTC and exchange-traded instruments as of last night. (The back office processing of today's positions is not completed until 22:00.) The MtF values in the draft Basis MtF Cube are computed using today's market data and yesterday's global scenarios. Yesterday's Basis MtF Cube is guaranteed to span yesterday's position. On a typical day, the characteristics of 5–10% of the OTC positions may change (e.g., a new reset rate for a swap or a floating rate note); the characteristics of the majority remain unchanged. By exploiting the additivity of MtF values, the Spadina Bank is able to start processing six hours before the end-of-day position is consolidated at 22:00. Several MtF engines are run in parallel, each processing a sub-set of the basis instruments across the relevant scenarios and time steps. The calculation of the draft Basis Cube is easily completed by 22:00 local time.

Update the Draft Basis MtF Cube at the End-of-Day: When the position files from the back office systems are available at 22:00, a process is run to

- identify the delta positions, the changed and new positions with respect to the positions of the previous day
- identify additional basis instruments required to map new or changed positions
- compute the MtF value of the additional basis instruments
- augment the Basis MtF Cube with these new values.

Each of these steps is a simple arithmetic operation; by 23:00 the Tokyo Basis MtF Cube generation is complete. This is the only time that the MtF values of these instruments are calculated; no new revaluations are required for global reporting or for multiple reporting hierarchies.

Generate Tokyo Interim Portfolio MtF Cube: The Basis MtF Cubes are mapped into Product MtF Cubes and finally combined into Portfolio MtF Cubes across market risk and credit risk limit structures.

Generate Tokyo risk reports: The interim Portfolio MtF Cube is post-processed to calculate the risk measures of interest, summarized in Table 1.3. The interim reports are available at 00:00, one hour after the interim Portfolio MtF Cube is available.

Prepare and distribute global risk scenarios	18:00	23:00	20:00	01:00
Generate validated local risk reports	08:00	15:00	09:00	14:00

Table 1.7: New York head office daily process

Correct Errors and Omissions: The daily process of validating position data at 07:00 local time is unchanged. However, in the distributed MtF framework the information from the errors and omissions report is used to update the interim Basis MtF Cube in a manner similar to the way end-of-day position data is used to update the draft Basis MtF Cube.

The ability to correct transaction errors without re-simulating the entire portfolio is a key advantage of persistent storage of the MtF Cubes. For example, only the MtF values of instruments affected by an exchange rate must be recomputed if the exchange rate is corrected. The Product and Portfolio MtF Cubes for only the corrected and omitted transactions are calculated and stored in the MtF database in Tokyo; the MtF Cubes of the erroneous positions are reversed.

Produce and Distribute Risk Reports: The risk reports are regenerated. Incrementally regenerating risk reports takes minutes rather than the many hours required to re-simulate the entire portfolio. Final, validated risk reports are delivered to the trading desk and senior management by 09:00 on day *T+1*.

London Branch Processing. The processing steps in the London branch, which are similar to those in Tokyo, begin at 16:30 and finish at 23:00. For the most part, the London Basis MtF Cube consists of basis instruments unique to the London portfolio. However, basis instruments for products traded continuously over 24 hours in multiple markets, for example the JPY/USD spot rate, are found in the Tokyo, the London and the New York Basis MtF Cubes. This eliminates the need to transmit Basis MtF Cubes; only Portfolio MtF Cubes are transmitted to the head office.

Interim draft reports are available at 00:00. By applying corrections at 07:00, the London risk manager produces validated local reports by 09:00 on day *T+1*. The London branch process is summarized in Table 1.5.

New York Branch Processing. The unvalidated New York Portfolio MtF Cube is generated by 00:00. Between 07:00 and 08:00, trade corrections are applied to the New York MtF Cube; the New York risk manager produces validated local reports by 09:00 on day *T+1*. The New York branch process is summarized in Table 1.6.

Upon completion of the daily process in New York, the head office of the Spadina Bank has a validated, distributed MtF database. Producing global risk reports is now straightforward.

Many of the features of a distributed MtF approach can be incorporated in the centralized MtF approach presented in Section 1.3. In fact, the distributed process illustrated in Figure 1.5 can be implemented in a centralized MtF environment. In a modified centralized MtF approach, draft position data is transmitted to the head office as soon as the market data is available and used to generate a draft Basis MtF Cube. Because the bulk of the processing is completed earlier in the overnight processing window, the bank can realize improved performance with respect to its reporting objectives. Subsequently, the branch offices must transmit to the head office both the end-of-day updates and the errors and omissions; the head office must use these data to update the Basis MtF Cube.

The relative merits of the two MtF approaches depend on the benefits associated with centralized versus distributed processing and the implications of the nature, quantity and frequency of the data transmissions associated with either approach. Therefore, the Spadina Bank can achieve its objectives by adopting either a decentralized process or a more complex, centralized process based on the MtF framework.

Head Office Processing

The New York head office is responsible for

- creating and distributing global scenarios
- generating and distributing global risk reports.

The head office daily process is summarized in Table 1.7.

Prepare and Distribute Stress Scenarios: At 19:00 global VaR, stress test and sensitivity scenarios for the next day are prepared centrally and then distributed to London, New York and Tokyo.

Generate Global Risk Reports: We consider the production of global VaR reports in some detail; similar post-processing produces the sensitivity, stress-testing, equity risk and credit risk reports.

VaR reports must be prepared across the market risk limits structure and the credit risk limits structure illustrated in Figures 1.1 and 1.2. The Global Portfolio MtF Cube is created by amalgamating the MtF Cubes of Tokyo, London and New York. The global VaR calculation is based on 12,000 MtF values (1,000 scenarios in each of the four

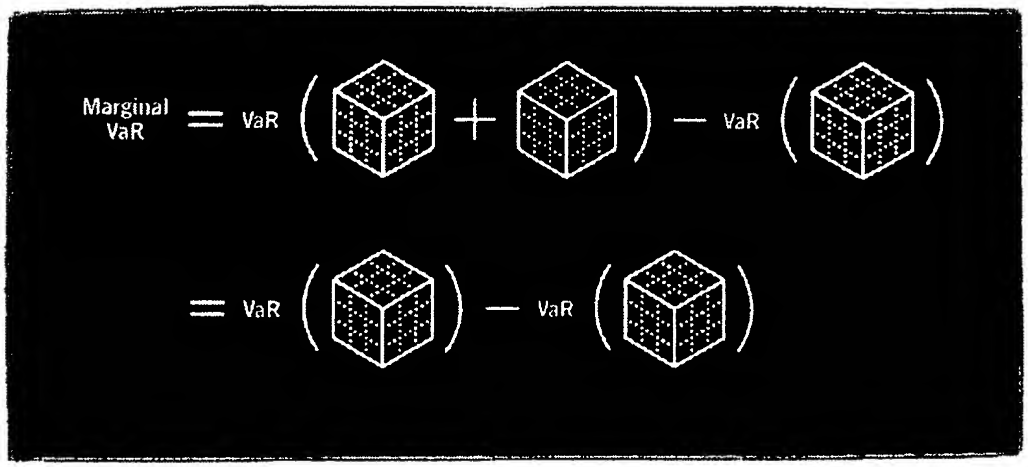

Figure 1.6: Marginal VaR calculation

scenario sets; each scenario has three time steps) for each portfolio; the local Basis and Position MtF Cubes are not required. The head office post-processor calculates global VaR based on the Global Portfolio MtF Cube.

To compute VaR measures across the credit risk limit structure, Tokyo, London and New York compute and return the MtF Cubes for each of the counterparty-product portfolios at the lowest level of this portfolio hierarchy.

The daily cycle of local and global reporting and generation of new stress scenarios is completed in time for the next day's operations in New York. Validated global reports are available in London at 14:00 and in Tokyo at 23:00 local time on day *T+1*. Moreover, the availability of the validated local Portfolio MtF Cube and the draft Global Portfolio MtF Cube at the start of trading on T+1 enables risk managers in the branch and head offices to perform marginal deal and *what-if* analyses.

An additional benefit arises from the implementation of the distributed process. Since draft reports are available for each of the branches by 05:00 GMT, the London branch can receive an interim Global Portfolio MtF Cube and interim global report before the start of the trading day in London on *T+1*. The Tokyo branch can receive an interim Global Portfolio MtF Cube and report by 14:00 on day *T+1*.

Marginal Analysis

The MtF approach streamlines the overnight process and facilitates the implementation of an enterprise risk system that feeds marginal pre-deal analysis information both to the local and head office decision makers. Recall from Part 2 that a key feature of MtF is that risk measures are calculated in the post-Cube stage. The local and Global Portfolio MtF Cubes are a platform upon which to analyze the incremental impact of deals considered throughout the day. Therefore, only the MtF values of proposed transactions need to be calculated. This enables the marginal deal application to be used in a near real-time environment.

The MtF values for a new deal calculated using the same scenarios, time steps and regimes used in the overnight process are stored in a Position MtF Cube. The difference between the VaR of the Global Portfolio MtF Cube consolidated with this Position MtF Cube and the original global VaR is the marginal VaR (Figure 1.6). The same principal can be used to calculate the marginal effect of a new trade on other risk measures of interest (e.g., counterparty exposure, capital, unrealized profit and loss).

In this section three simple examples illustrate how MtF values are used to compute risk measures in a MtF framework. The examples are based upon a marginal VaR

New York options desk	5	5	8	8
New York derivatives department	20	25	22	25
New York branch	40	50	38	50
Global	70	80	69	80

Table 1.8: Impact of increasing VaR limits on VaR (millions USD)

Original 95% total exposure	5.0	8.0
95% total exposure with proposed deal	4.8	8.2
Marginal 95% total exposure	-0.2	0.2

Table 1.9: Comparison of counterparty total exposure

application and include

- evaluation of a request for additional limit capacity
- pre-deal limit check for VaR
- marginal deal analysis on counterparty credit exposure.

Request for Additional Limit Capacity

The Spadina Bank's market risk committee meets weekly to review its market risk exposure, both globally and at the desk level, and to evaluate limits requests. Today, the New York options desk has requested a 60% increase in its VaR limits, from 5 million USD to 8 million USD, in response to increased business opportunities. The desk head believes that the composition of the business will remain more or less unchanged with the increased volume. During the meeting, the senior quantitative analyst uses the marginal deal application to scale the Portfolio MtF Cube of the options desk portfolio until its VaR reaches 8 million USD. This corresponds to a 10% increase in business volume. He then re-computes the VaR of all higher level portfolios. The result is presented in Table 1.8.

Based on this analysis, it is apparent that the New York options desk is acting as a natural hedge for the New York branch and the Global portfolio because the New York and global VaR are actually reduced. The increase in VaR limits for the New York options desk is approved. Limits for New York derivatives, the New York branch and the global holdings remain unchanged.

Pre-Deal Limit Check

A bond trader in Tokyo is within 20,000 USD of her VaR limits. Prior to purchasing a new corporate bond she enters the bond into the marginal deal application. The MtF of the proposed, new transaction is generated using the same scenarios, time steps and regimes used for the overnight calculation. The MtF of the bond is then combined with the MtF of her book and the VaR at 95% is recalculated.

The new trade increases the VaR of the portfolio by only 10,000 USD. The limit is not exceeded and the trade is executed. The entire process takes less than a minute to complete.

Marginal Deal Analysis on Counterparty Credit Exposure

A trader on the interbank desk is responsible for hedging the risk from the swaps desk in the interbank market. He receives the same quote of LIBOR + 20 bp from two large trading counterparties, Richmond and Manor. The marginal VaR and the return of each deal are the same.

Using the marginal deal application the trader enters the proposed transaction into the system twice, once with Counterparty Richmond and once with Counterparty Manor. He then instructs the system to compute the marginal credit exposure of each deal. The marginal credit exposure of each deal is the difference in counterparty total exposure with and without the proposed deal. (Total exposure is defined in Part 2, Chapter 7.)The results of this analysis are presented in Table 1.9. The netting provisions in this new deal actually reduce the counterparty total exposure to Manor. Based on this analysis, the transaction is executed with Manor.

In this chapter, the benefits of a distributed process built on the MtF framework have been presented. Distributing the risk calculation, pipelining and parallelizing the process, then aggregating the MtF Cubes enables the production and distribution of global risk reports in a timely manner. The architecture required to support these distributed operations is presented in the next chapter.

CA 02290888 1999-11-26

Distributed Mark-to-Future Architecture.

AS DISCUSSED IN CHAPTER [illegible], THE CHIEF PROBLEM WITH A CENTRALIZED ARCHITECTURE IS THE NON-ADDITIVITY OF THE KEY ENTERPRISE RISK MEASURES. THOUGH THE ENTERPRISE MEASURES ARE NON-ADDITIVE, MARK-TO-FUTURE CUBES ARE ADDITIVE (ASSUMING THE SCENARIOS AND TIME STEPS ARE CONSISTENT), AND THAT MAKES ALL THE DIFFERENCE. THIS CHAPTER PRESENTS THE ARCHITECTURAL CONSIDERATIONS UNDERLYING AN ENTERPRISE RISK MANAGEMENT SYSTEM BASED ON THE MARK-TO-FUTURE FRAMEWORK.

A MtF Cube consists of the value of an individual basis instrument, product or portfolio of positions, under various scenarios and through time. Therefore, each element in a MtF Cube is a MtF value for an instrument, product or portfolio at a given future time under a given scenario. If we obtain two such MtF Cubes, one for product A and another for product B, then the MtF Cube for a portfolio comprising one unit each of A and B is the element-wise sum of the Product MtF Cubes of A and B. The element-wise sum is obtained by adding the MtF value for product A under scenario j at time t to the MtF value for product B under scenario j at time t, as shown in Figure 2.1, for each scenario and each time step. This is true because the value of a portfolio at a certain point in time under a certain scenario is indeed the sum of the values of its constituent positions at that same point in time and under that same scenario.

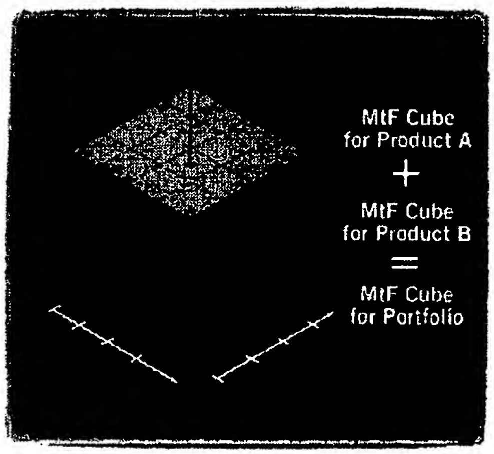

Figure 2.1: Element-wise additivity of MtF Cubes

The implications are twofold. First, there is a performance advantage to be gained. The Basis MtF Cube need only be computed once, and can subsequently be combined and recombined in various ways without performing long re-computations to yield different portfolio structures and different risk measures. The second implication is that it becomes possible to apply a decentralized, distributed architecture to the enterprise-wide risk management system.

The system is distributed because it becomes possible to distribute the computations of MtF Cubes throughout the enterprise. As long as MtF engines scattered throughout the enterprise are provided with a consistent set of scenarios, each MtF engine can independently compute its MtF Cubes. This high degree of distribution yields a scalable solution, one that efficiently shares the computing resources spread around an enterprise and is simpler to implement, as each distinct business unit can configure and operate its own MtF engine.

The MtF Cubes computed locally can be consolidated centrally and the desired enterprise risk measures computed. As well, enterprise Portfolio MtF Cubes can be distributed to the business units where they can be joined with locally-computed Portfolio MtF Cubes to compute marginal contributions efficiently and in a decentralized fashion.

Penny (1999) reviews the architectural considerations in a distributed enterprise risk management system that exploits the MtF paradigm. This chapter outlines the high-level architecture of a MtF implementation, describes key components in more detail and summarizes the technology benefits of implementing the MtF framework.

High-Level Architecture

The practical implications of a distributed MtF process are challenging from an architectural standpoint. Generally speaking, a distributed system architected around the MtF framework consists of the following architectural components, summarized below and illustrated in Figures 2.2 and 2.3.

1. An **input data architecture** for holding configuration information and for collecting and distributing terms and conditions data, position data, credit data, raw market data and historical time series data.
2. A **curve engine** that takes raw market data and calibrates the curves used to drive the instrument pricing functions.
3. A **scenario engine** for producing scenarios from time series data and *ad hoc* sources.
4. A **MtF engine** that is fed curves, scenarios, a set of basis instruments and a set of time points and produces the MtF results.
5. A **distributor component** that allocates scenarios and basis instruments to MtF engines.
6. A **MtF database** capable of holding the results of MtF computations, an access layer and an indexing facility that points to what is stored where.
7. A **hierarchy server** that describes how positions combine within the enterprise into portfolios of holdings.
8. A recursive procedure called a **MtF post-processor** that combines the MtF values according to information in the hierarchy server and computes various statistics on the future distribution to yield enterprise risk measures.
9. A collection of MtF post-processors organized in some functional way and known as a **MtF aggregation engine** aggregates results from the MtF post-processors and synchronizes results among MtF aggregation engines.
10. The **front-end applications** required to configure the system and display final results.

In a distributed architecture each of these components is replicated, meaning that many identical copies of these engines run throughout the enterprise, cooperating with one another to solve the enterprise problem.

The high degree of distribution is made possible because the problem can be disaggregated along the scenario dimension, the instrument dimension, the time dimension and the portfolio dimension. Referring to the logical view illustrated in Figure 2.2, each MtF engine values the subset of basis instruments and scenarios it is allocated. The MtF database can support the distribution of data across the same instrument and scenario dimensions and across the time dimension. The MtF post-processor disaggregates the problem along the portfolio hierarchy dimension. That is, an individual MtF post-processor performs the computations across each node in a portfolio hierarchy. This distribution is needed to deal with the potentially massive data flow.

While all of this data must inevitably be loaded and accessed, it does not necessarily need to pass over the network. Referring to the physical view illustrated in Figure 2.3, it is possible to group a MtF engine, a part of the MtF database and a MtF post-processor on the same physical machine. This grouping allows each physical processor to transfer data in parallel with the others. This network parallelism is the key to the MtF architecture.

The remainder of this chapter deals with three particularly important components of the architecture: the MtF engine for computing MtF Cubes, the MtF database for holding these Cubes and the MtF aggregation engine for post-processing MtF Cubes to yield useful enterprise measures.

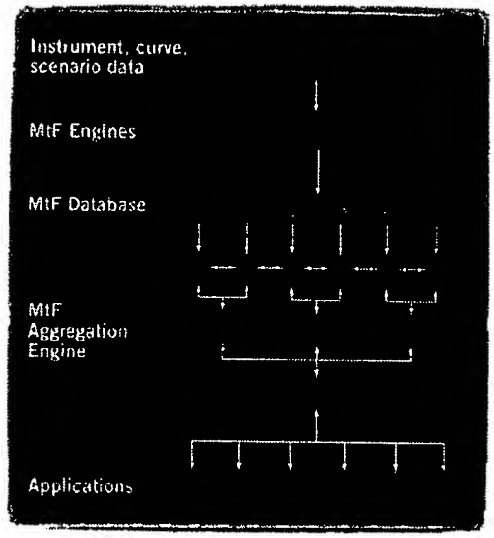

Figure 2.2: Logical view of MtF architecture

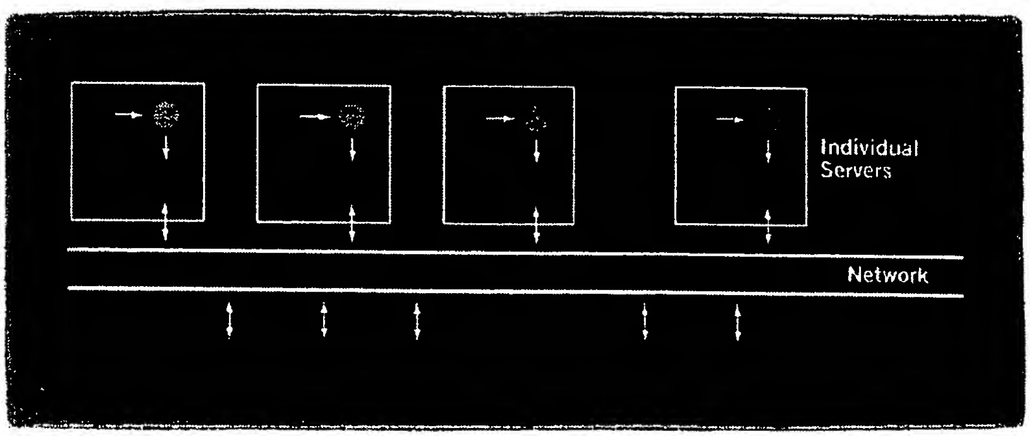

Figure 2.3: Physical view of MtF architecture

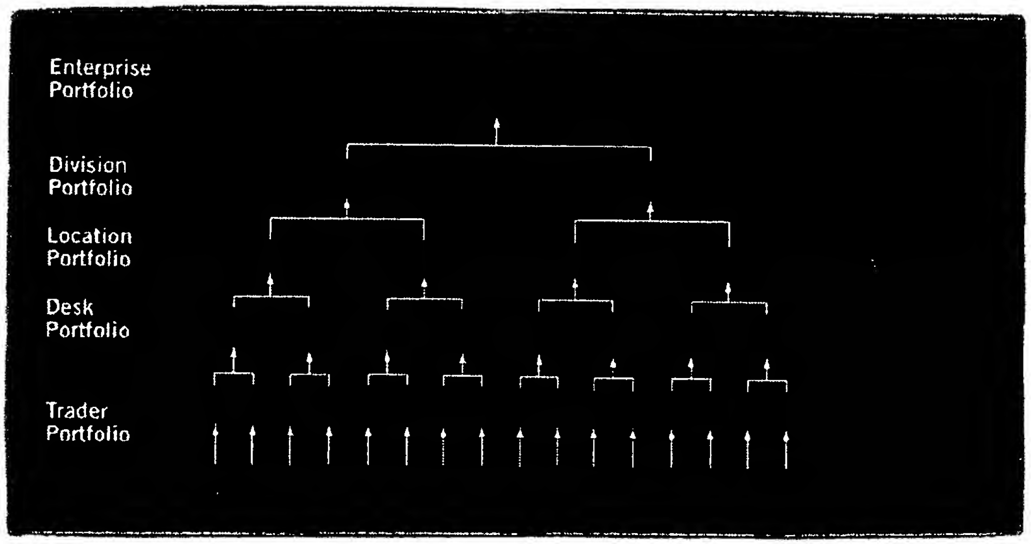

Figure 2.6 MtF post-processing for an organizational hierarchy

Mark-to-Future Engine

A vital component of the system is a server capable of computing the MtF Cubes and storing them in a database. Each engine should have access to the static data environment required to perform the computations. This includes the detailed terms and conditions of financial products, the curves necessary to perform a mark-to-market valuation and the scenarios that are used.

An individual MtF engine fields requests to MtF basis instruments under scenarios. The MtF engine loads the necessary static data and produces the required MtF computations under a pre-configured set of time steps. A distributor component divides any problem requiring a MtF solution into many such invocations of the MtF engine, effectively distributing the computation. Distribution leverages one of the key advantages of the MtF framework—the fact that the calculations are trivial and can be performed in parallel along both the instrument and the scenario dimensions.

All financial choices in a MtF engine should be configurable. In particular, it should be possible to alter the functionality of the MtF engine using dynamic load technology in order to accommodate new financial products. For example, of prime importance in the MtF engine is the ability to produce an exceptionally accurate MtF value in short order, but it is also useful to be able to trade that accuracy for improved computation time, either through instrument consolidation or through faster, though less accurate, pricing approaches.

Mark-to-Future Database

As described above, a global financial institution must store large amounts of MtF data. Only by spreading this data across many data storage devices can acceptable data flow into and out of the MtF database take place. Therefore, the prime driver for the design of a MtF database is the need to store large volumes of data and to populate and access that data in a distributed fashion.

Basis MtF Cubes computed by a given MtF engine should reside on their own data storage device near that of the MtF engine. A distributed indexing system keeps track of which Basis MtF Cubes are stored where. An access layer on top of the whole database system provides a virtual view of any combination of the individual databases. Such a system is sometimes referred to as a **virtual database**.

Mark-to-Future Aggregation Engine

The key objectives of post-processing are to combine Basis MtF Cubes into Product and Portfolio MtF Cubes and to compute various statistics on them. In the sense used here, a portfolio refers to any hierarchical grouping of instrument holdings. For example, a bank can divide its total instrument holdings by a static, organizational portfolio hierarchy (e.g., trader, desk, location, division

and enterprise). Alternatively, the instrument holdings can be divided along more dynamic lines, such as by instrument type, currency, maturity date and so on. Each of these alternative portfolio views can provide insight into the distribution of risk among the total holdings.

In general, at the base of any portfolio hierarchy are the leaf portfolios, those that contain positions in financial products. Portfolios above the leaf level contain only higher level portfolios. In the MtF framework, associated with each leaf portfolio is a part of the distributed database of Basis MtF Cubes or a reference to the relevant part of the distributed database, and associated with that is an instance of a MtF post-processor. Intermediate-level MtF post-processors receive the Portfolio MtF Cubes for the lower level portfolios from the lower level post-processors. These intermediate Portfolio MtF Cubes are then forwarded to yet higher level MtF post-processors, and so on up the chain until they reach a MtF aggregation engine.

Figure 2.4 illustrates the *roll-up* processing in an organizational portfolio hierarchy. The leaf portfolios correspond to trader portfolios; associated with each trader portfolio is a MtF post-processor and a part of a MtF database. Intermediate-level MtF post-processors are associated with the desk and location portfolios in the intermediate rungs of the hierarchy. A MtF aggregation engine is associated with each division. The output of the aggregation engines is synchronized to produce MtF values for the enterprise.

The Portfolio MtF Cube associated with a trader portfolio is forwarded to the MtF post-processor for the appropriate desk portfolio. There it is simply summed with the Portfolio MtF Cubes of other traders on that desk (Figure 2.5). The Portfolio MtF Cube associated with the desk is forwarded to the MtF post-processor for the appropriate location, where it is summed with the Portfolio MtF Cubes of other desks in that location. The Portfolio MtF Cube of each division is constructed in the same manner. Finally, the Portfolio MtF Cubes of each division are summed by the associated aggregation engine and synchronized with one another to produce the enterprise Portfolio MtF Cube.

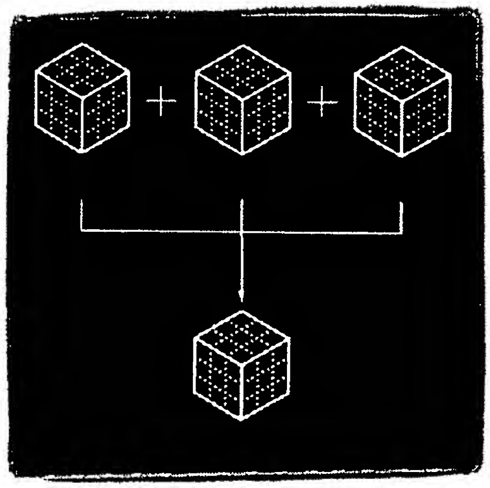

Figure 2.5: Post-processing the Portfolio MtF Cube for a desk

Accordingly, only aggregated Portfolio MtF Cubes are ever propagated across the network, vastly reducing the data transfer requirements needed to solve the enterprise problem.

To minimize network traffic, a clever allocation of basis instruments to machines is necessary. In particular, those basis instruments referenced from commonly requested leaf-level portfolios should be grouped whenever possible. This maximizes the chances that only Portfolio MtF Cubes must be transmitted. Many potential portfolio structures exist, implying different leaf-level portfolio compositions for each. This suggests a time/space trade-off. To minimize network traffic and hence maximize performance, it is desirable to replicate the Basis MtF Cubes for a given instrument across more than one MtF database. In this way, a given MtF database is more likely to contain all the data required to assemble a leaf-level portfolio.

Each MtF aggregation engine can field a request for any portfolio. If the data to satisfy that request is not available locally, the aggregation engine forwards the request to another aggregation engine that has the relevant data. This works top-down, as described previously, but can also work bottom-up in the sense that a local MtF aggregation engine can be asked to combine the MtF Cube of a newly proposed *what-if* deal with the Portfolio MtF Cubes of higher level portfolios. This yields the all important marginal analysis numbers capable of driving efficient capital allocation within a financial institution.

Benefits of a Distributed Mark-to-Future System

The benefits of a distributed MtF architecture are summarized below.

Parallelize and pipeline processing. Since MtF values are additive by scenario or by instrument, different portions of the global portfolio can be processed as soon as the instrument, curve and scenario data for those portfolios are available. This is known as pipelining. Because MtF values are additive, different subsets of the portfolio and scenarios can be computed in parallel. This is known as parallel processing. Together, they reduce the total time required for processing. In some cases, it may be efficient to duplicate basis instruments in different Basis MtF Cubes; for example, a JPY/USD spot instrument that is traded in more than one location in a 24-hour cycle (as described in Chapter 1). In other cases, it may be prudent to distribute the calculations. In the limit, each value in a MtF Cube can be computed on a different machine. In practice, the degree of duplication versus distribution is chosen to optimize system performance and to satisfy local and global risk reporting requirements.

Fast error correction. Incorrectly entered or missing positions are inevitable; these errors need not invalidate the entire day's risk results. The Basis MtF Cubes computed from incorrectly entered positions can be updated with the Basis MtF Cubes of the correct positions since regenerating risk measures is a simple matter of post-processing the MtF result.

Integrated systems architecture for market, liquidity and credit risk. The MtF methodology and the MtF architecture that supports market, liquidity and credit risk measurement and management is based on common system components. The same MtF engines, scenario engines and aggregation engines are used for market, liquidity and credit risk. The required risk measures are obtained through proper choices of scenarios, portfolio evolution and post-processing of the MtF result.

This chapter describes how a bona fide, leading financial institution successfully implemented the MtF framework to measure and manage its risks more effectively and to allocate capital, positioning itself ahead of its industry counterparts in the field of risk management. It shows how the MtF framework is tailored to meet the needs of the financial institution and describes the assumptions made during implementation.

Using the MtF framework, HypoVereinsbank, a leading German financial institution, can produce a number of different risk measures and decompose risk at different levels in the portfolio in a timely manner. This chapter is organized as follows. After introducing HypoVereinsbank and describing its business objectives, the six steps of the MtF methodology (outlined in Part 2) as adapted to the needs of HypoVereinsbank, are presented. The details of the implementation and the architecture are presented, the MtF implementation is benchmarked against a standard implementation and the benefits to the bank of implementing the MtF framework are summarized.

Mark-to-Future

About HypoVereinsbank

Formed in September 1998 by the merger of two Bavarian banks—Bayerische Hypotheken und Wechseln Bank and Bayerische Vereinsbank—HypoVereinsbank has 40,000 employees. With roughly 450 to 500 billion Euro in assets, it is one of the top five European banks. HypoVereinsbank is one of the largest real estate finance companies in the world and also has a substantial focus on corporate banking, asset management, retail and selected investment banking activities.

HypoVereinsbank distinguishes itself from other large, global European banks by deploying its core competencies—real estate financing and corporate banking—in selected regions worldwide. The bank's trading business, while not the center of the bank's activities, is still a sizeable component, contributing roughly one quarter of the bank's pre-tax profits.

The bank believes that to inspire client confidence and appease regulators, sophisticated risk management technologies are required. HypoVereinsbank's goal is to build a measurement system that enables it to achieve group-wide consistency of methodology and transparency of results, as well as functional integration of various classes of risk and P&L. In addition, it wants to be able to use a number of different risk measurement methods including both historical and Monte Carlo simulations and to decompose risk at different levels in the portfolio. Finally, selected information and analytics must be delivered in a timely fashion to those who can act upon the information.

To meet these objectives, HypoVereinsbank has implemented a risk management system based on the MtF framework using RiskWatch (Algorithmics 1999b) as the MtF engine. To date, the bank has completed the mapping into RiskWatch of all interest rate products in its trading center in Munich and plans to map products in its other major trading centers (in London, New York and Singapore) in stages, within the year. The current status should be viewed as a first step in a progressive implementation of the MtF framework.

Six Steps in Marking-to-Future at HypoVereinsbank

The risk management process in HypoVereinsbank is based on the six steps of the MtF methodology described in Part 2:

1. Define the basis instruments.
2. Define the scenarios and time steps.
3. Simulate the instruments to generate a Basis MtF Cube.
4. Map the financial products onto the Basis MtF Cube.
5. Map the portfolio regimes onto the Product MtF Cube.
6. Post-process to produce risk and reward measures.

This section describes how HypoVereinsbank adopted the six steps in the MtF methodology to implement an inclusive and unified framework to measure the interest rate market risk of approximately 100,000 positions in its trading book daily. Instruments include exchange-traded and over-the-counter (OTC) interest rate derivatives and their underlying instruments. Cash flows from the banking book are included to calculate enterprise-wide interest rate risk and for asset and liability management.

Because this is the initial phase of a full-scale MtF implementation and because of the short-term horizon and the predominance of OTC positions in the book, the steps outlined in Part 2 are amended and abbreviated to suit the implementation at HypoVereinsbank. In future stages, the MtF implementation will be extended to include both market and credit risk management at all of the major trading centres. Longer time horizons, multiple time steps and dynamic portfolio strategies will be included to account for credit risk. At that stage the performance improvements that can be realized by using basis instrument mapping will be significant.

Step 1: Define the Basis Instruments

The bank uses information at the product level, rather than at the instrument level, because the majority of its positions are OTC. Therefore, the mapping of instruments to products is one-to-one. Position data are obtained from the bank's Treasury Database, an internally developed solution that houses worldwide position data.

Step 2: Define the Scenarios and Time Steps

The current implementation at HypoVereinsbank focuses on market risk for its trading businesses. As such, a relatively short time horizon is appropriate. HypoVereinsbank measures risk over both a one-day and a 10-day time horizon, using a single time step in each case. This satisfies both internal risk management and regulatory requirements.

The use of a 10-day holding period provides the bank with a suitable measure for risk for most of its liquid, tradable instruments. Measures based on this horizon are often used to measure and allocate economic capital, to manage market risk limits and for portfolio management on the trading desks. It also allows the bank to meet Bank for International Settlement (BIS) regulatory requirements (BIS 1996). The bank uses a one-day holding period for backtesting purposes.

In MtF, risk depends on possible future events described by scenarios. Scenarios can be based on history or subjective estimates or derived from models. Atypical events captured by stress scenarios include the breakdown of normal correlations, discontinuous price moves and atypical market shocks. One of the more significant implications of the MtF framework is the ability to make explicit scenario choices, in particular with respect to stress testing, and to combine stress scenarios in a consistent manner with scenarios describing typical markets.

HypoVereinsbank uses a number of different scenario generation techniques to describe future events. Monte Carlo scenarios based on historical data are used to provide forecasts for typical events. Stress test scenarios generated by individual risk managers are used to provide subjective estimates of the future and to account for atypical events.

Monte Carlo Scenarios: HypoVereinsbank generates approximately 4,000 single-step Monte Carlo scenarios in RiskWatch using a variance-covariance (VCV) matrix in RiskMetrics format (Longerstaey and Zangari 1996) obtained from Olsen & Associates. The VCV matrix is created using 250 days of historical data and 550 risk factors including interest rates, FX rates and implied volatilities for G12 countries and emerging markets. This method of scenario generation is based on explicit assumptions (e.g., one time step, no mean reversion and a joint-normal distribution of factor log returns).

Monte Carlo scenarios provide HypoVereinsbank with a view of the future, using typical circumstances from the past as a guide. These scenarios account for past events that are not based solely on historical observations. This allows the bank to overcome many of the data issues commonly associated with the use of historical simulation. The bank chooses to use 4,000 scenarios because this number provides a good balance between performance and accuracy—the number is sufficiently large to allow for a high degree of confidence in risk measurement without adversely affecting the timeliness of risk reports.

Stress Test Scenarios: The use of Monte Carlo scenarios allows the bank to measure risk under normal or typical circumstances; however, it does not account for the extreme market moves or atypical events which have occurred since the beginning of organized markets. Effective risk management must minimize surprises, and hence must cover atypical events. These events include the breakdown of normal correlations, discontinuous price moves, sudden decreases in liquidity and atypical market shocks.

To account for these situations, HypoVereinsbank has implemented daily stress testing procedures to meet the needs of the enterprise risk management group. Individual risk managers create their own scenarios, based on their subjective views of the future. These stress tests may include shocks to individual risk factors or node points, shocks to driver curves or spread curves and shocks to FX rates and volatilities.

The stress testing procedure ensures that the bank accounts for atypical events through subjective views of the future and, as such, minimizes surprises in earnings. User-defined scenarios are an essential ingredient in effective risk management because they allow the bank to accommodate an array of opinions about the future. It is relatively easy for risk managers to assess the reasonability of these scenarios because they can be viewed graphically. Scenarios are descriptive and therefore easy for non-technical managers to define and understand.

Step 3: Simulate the Instruments to Generate a Basis MtF Cube

HypoVereinsbank calculates the mark-to-market value of its holdings in RiskWatch using closing prices, zero curves and volatilities obtained from its trading systems nightly. Static data and terms and conditions data are obtained from yet another database (Gips). All input data is fed into the bank's input database (Input DB) and then subsequently mapped into RiskWatch, where it is marked-to-market using both Algorithmics' and custom-developed, proprietary pricing models.

The MtF value of each position at the horizon, under each scenario in both the Monte Carlo and stress test scenario sets is computed, accounting for the events that occur between the horizons and their impact on P&L.

Step 4: Map the Financial Products Onto the Basis MtF Cube

Given that the majority of the positions are OTC, the mapping of instruments to products is one-to-one. Therefore, additional mapping to products is not required.

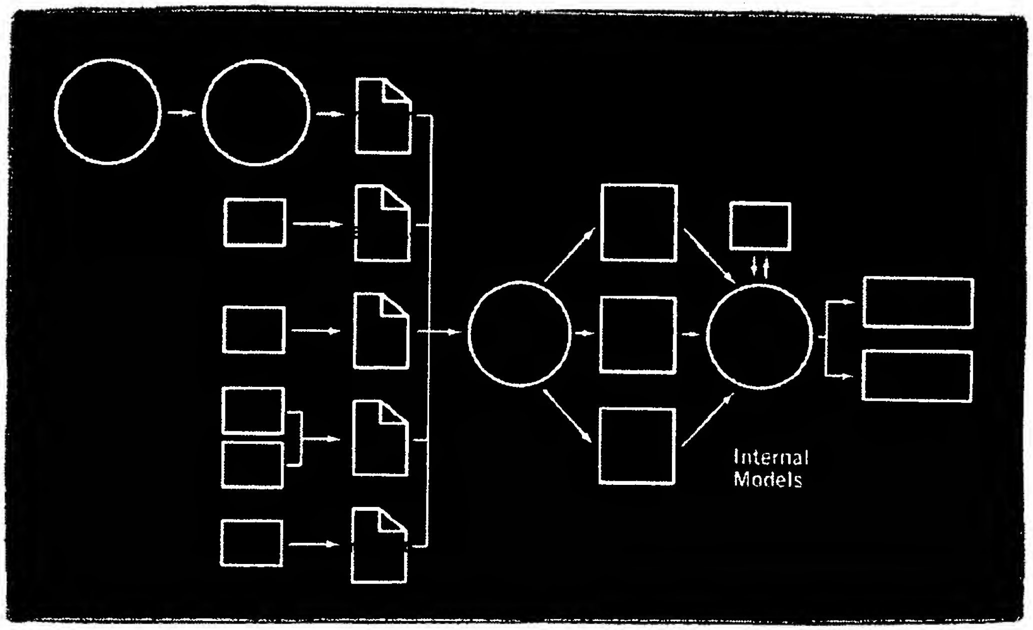

Figure 3.1: HypoVereinsbank's high-level process

Step 5: Map the Portfolio Regimes Onto the Product MtF Cube

Given HypoVereinsbank's fairly short time horizon for market risk, it is not necessary to consider changes to the composition of the portfolios. The position MtF values by scenario are exported to the Output Database. Database stored procedures aggregate the position level results along portfolio lines.

Step 6: Post-Process to Produce Risk and Reward Measures

Once the steps above have been completed, the information needed to compute true forward profit and loss is available. These P&L values (the MtF value less the mark-to-market value) are used to compute VaR(99%), VaR contribution and maximum loss among other portfolio statistics.

HypoVereinsbank's high-level process is shown in Figure 3.1.

The sections that follow describe the high-level architecture deployed at HypoVereinsbank and provide details of the construction of the bank's MtF Cube.

HypoVereinsbank Architecture

The MtF framework has significant technological implications. Using this framework, financial institutions can decentralize risk calculations and realize performance gains, which result in vast improvements in total computing time. HypoVereinsbank has implemented the MtF framework using the architectural components shown in Figure 3.2 and summarized in Table 3.1.

Building the Portfolio MtF Cube

This section describes how HypoVereinsbank computes one-period MtF values for each position under each scenario and time horizon. It explains how the bank builds its Portfolio MtF Cube and demonstrates the benefits realized as a result of this approach.

HypoVereinsbank takes advantage of the additivity of MtF values by building the Portfolio MtF Cube in stages. Because MtF values are additive, the bank creates the MtF values for the components of a portfolio separately and then adds them together at a later stage

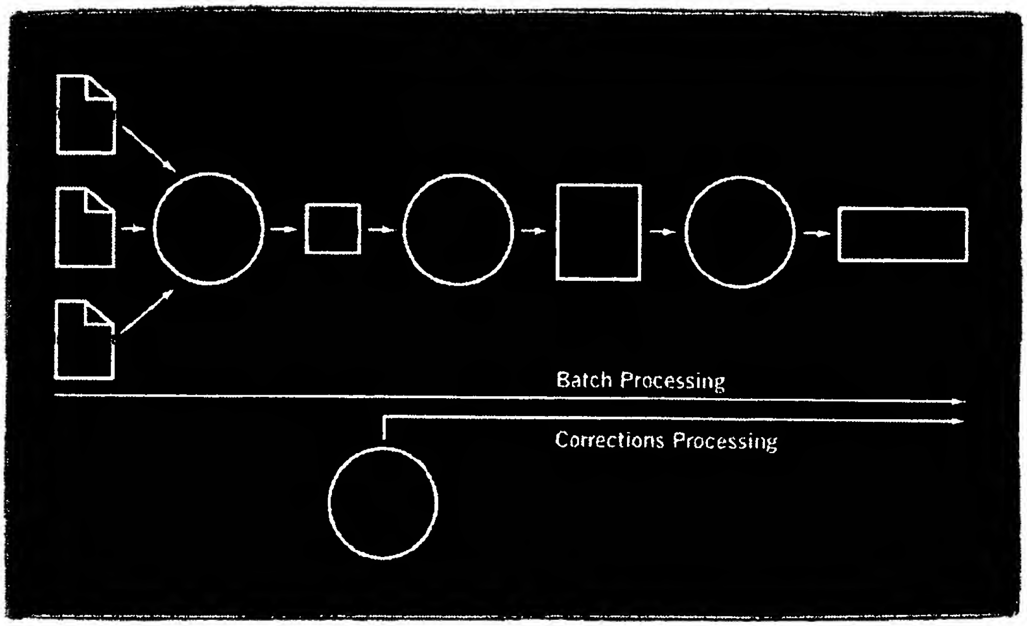

Figure 3.2: HypoVereinsbank's high-level architecture

1. Input Data Architecture	· Positions from trading systems (Intas and Kondor+) are stored in [illegible] DB and transferred to Input DB, then mapped into RiskWatch. · Market data and curves from the trading systems are fed to the Market Data database then to Input DB and mapped into RiskWatch. · VCV data from Olsen & Associates are fed into the Input DB then mapped into RiskWatch. · Limits from in-house limits system are fed into the Output DB. · Static and terms & conditions data from the in-house [illegible] DB are fed to Input DB and then to RiskWatch.
2. Curve Engine	· Trading systems (Intas and Kondor+) provide curve data.
3. Scenario Engine	· Monte Carlo scenario generation in RiskWatch uses VCV data from Olsen & Associates.
4. MtF Engine	· MtF valuation in RiskWatch is performed on a Sun Microsystems E10,000, running 8 RiskWatch sessions on 4 CPU each, with 4 GB per session.
5. Distributor Component	· Eight parallel RiskWatch sessions are managed on an E10,000, using 1 CPU each, with 2GB memory. RiskWatch sessions are managed using RiskScript.
6. MtF Database	· The Output DB is a 2 GB Oracle database on the E10,000.
7. Hierarchy Server	· Portfolio hierarchy is maintained and managed manually using a GUI to Input DB/Output DB.
8. MtF Post-Processor	· Lightweight applications/GUIs built in Visual Basic include interactive and batch reporting tools, which can aggregate VaR at different levels and provide a VaR contribution for each instrument and/or portfolio.
9. Front-End GUIs	· User-defined scenarios and stress tests as well as error correction for market data, position data, static data and systems errors via GUIs. · Visual Basic applications including Risk Viewer, NPV and [illegible] Viewer, Scenario Viewer and Stress-test Viewer.

Table 3.1: HypoVereinsbank's high-level architecture

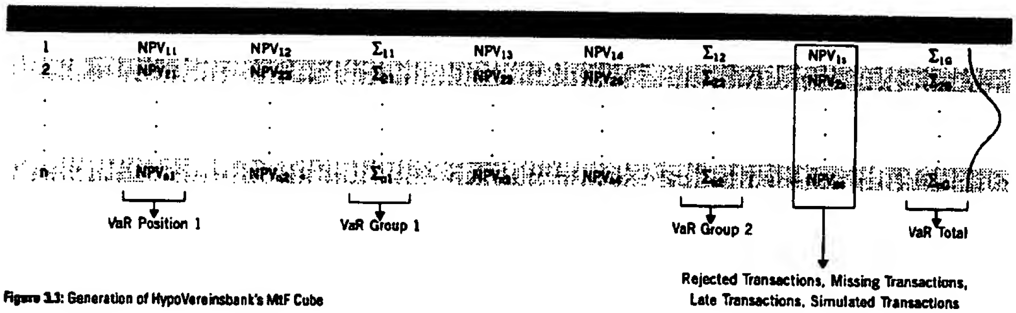

Figure 3.3: Generation of HypoVereinsbank's MtF Cube

to obtain portfolio risk measures. This results in improvements in both performance and accuracy.

First, the bank creates MtF values for its **draft positions**, closing positions from the previous day. This begins as soon as closing market data is available. Then, once the end-of-day positions have been processed, it creates MtF values for the changes in position. In the morning, after a review of the risk results and error reports, MtF values are obtained for any required corrections. Since only a small number of computations are required for error correction, this process is accomplished very quickly. Additivity allows the bank to begin its processing before final positions are available and to make corrections to MtF values without re-computing the entire portfolio. As a result, the bank obtains accurate risk measures in a timely manner.

Marking the Portfolio to Market

As soon as closing market data and curves are available from the trading systems, HypoVereinsbank begins the MtF process. To mark the portfolio to market, draft positions are fed from the bank's global position database (Treasury DB) into the risk management group's input database (Input DB). These are both in-house-developed databases. The positions are then mapped from the Input DB into RiskWatch. Current closing market prices, curves and volatilities are obtained from the bank's trading systems (Intas and Kondor+) and stored in a market data database. The market data is then fed into RiskWatch from the Input DB and used to mark the draft positions to market in RiskWatch using both Algorithmics models and proprietary pricing models developed using Risk++ (Algorithmics 1999a). Mark-to-market values are stored as the nominal scenario values in HypoVereinsbank's in-house-developed database, Output DB, which corresponds to the MtF database described in Chapter 2.

Creating the MtF Cube for the Draft Positions

HypoVereinsbank creates its Basis MtF Cubes for the draft positions using eight multiple, parallel sessions of RiskWatch. RiskWatch sessions are arranged by portfolio and by the number of positions in each portfolio, resulting in an even distribution of computations across RiskWatch sessions. The variance-covariance matrix obtained from Olsen is mapped into each RiskWatch session through the Input DB, as are the draft positions that have been allocated to each particular session. Each RiskWatch session uses the same VCV matrix and Monte Carlo seed to generate the same 4,000 single-step Monte Carlo scenarios for a 10-day time horizon. Then, each RiskWatch session values each draft position in that session, under each scenario. These MtF values are written to the Output DB and are stored in a series of arrays by position, scenario and time horizon (Figure 3.3) to form the Position MtF Cube. The process is repeated for the one-day time horizon. The entire process spans approximately two to three hours.

Adding the New Positions to the Inventory

Once the current day's positions are available from the trading systems, the changes in position from the previous day, the **delta positions**, are found by comparing yesterday's positions to today's positions. The same process as described previously is used to mark the delta positions to market and to future, using the same scenarios. This new MtF Cube is used to update the Position MtF Cube for the original inventory of positions in the Output DB and is later used to obtain the required portfolio risk measures.

Correction of Errors and Omissions

After this overnight process has been completed, HypoVereinsbank completes its rigorous error checking and control procedures. Automated error messages produced during the night indicate any system failures and any missing market or position data. Risk managers review these error reports as well as the preliminary risk measures to ensure that results are reasonable. If any

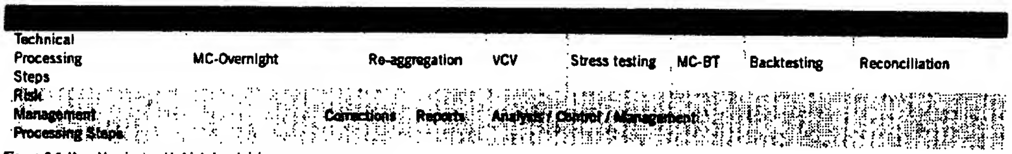

Figure 3.4: HypoVereinsbank's high-level risk management process

errors are found or if risk measures look unreasonable, risk managers are able to make corrections to the input data. For example, if an error is found in a position size, only the MtF values for that position are re-computed. If the DEM/USD FX spot rate is incorrect, only the MtF values for positions affected by the DEM/USD FX spot rate are re-computed. After the corrections have been made and verified, an application is used to re-compute MtF values only for the positions that have been affected. These MtF values are generated using the same scenarios and procedures described above. The overnight process is summarized in Figure 3.4.

A Universe of Risk Measures

MtF values are the essential ingredient in all risk measures. Once these values have been obtained, the relevant measures of risk may vary and are context dependent. With MtF it is now possible to move away from the traditional paradigm where a complex and monolithic application must be all things to all people. Instead, lightweight applications with low development overhead can be created to meet the needs of individuals and business units within a financial institution.

HypoVereinsbank has taken advantage of this flexibility. Several applications built in-house enable risk managers to explore the input data and the results of the MtF calculations.

Once the MtF values have been stored in the Output DB, the bank uses an Oracle application developed in-house to calculate a number of risk measures including the P&L for each position under each scenario and time horizon. The application uses these P&L values to compute various portfolio statistics such as VaR(99%), VaR contribution, maximum loss, etc., at each level within the portfolio, from the highest level all the way down to the position level. The Risk Viewer application shown in Figure 3.5 can be used interactively to aggregate risk by portfolio, by currency, by instrument type, etc. at all levels in the portfolio hierarchy.

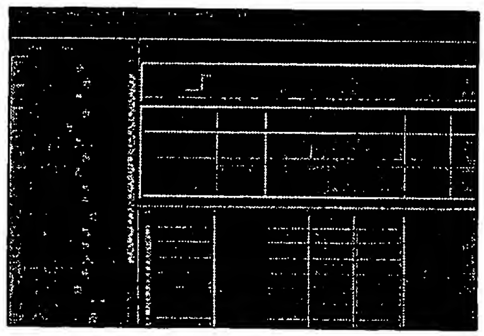

Figure 3.5: Risk Viewer

At each level within the portfolio hierarchy, the user can view the mark-to-market, the net present value (NPV) or MtF value under each scenario and the P&L under each scenario using the NPV and P&L Viewer application depicted in Figure 3.6.

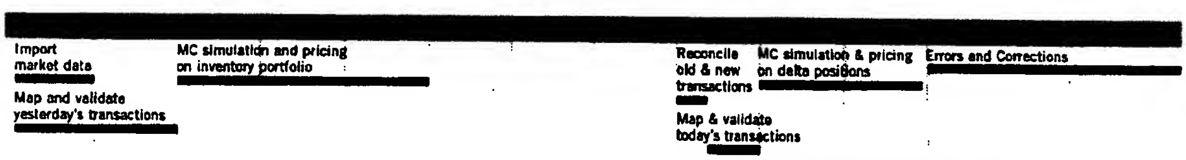

Figure 3.6: HypoVereinsbank's detailed daily process

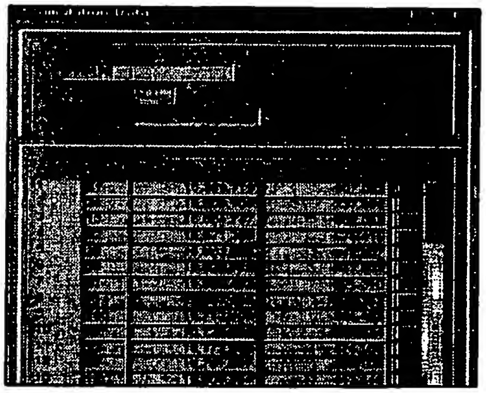

Figure 3.6: NPV and P&L Viewer

Another significant advantage of the Mark-to-Future framework is that risk managers can identify and examine the particular scenario that produces a given risk result. Using the Stress Test Viewer application shown in Figure 3.7, HypoVereinsbank risk managers can look at various risk measures such as the VaR(99%) or worst-case P&L. If they want to find out more about the global VaR scenario, for example, they click on an icon that takes them to the Scenario Viewer application shown in Figure 3.8. There they find a verbal description and a graphical representation of that scenario. This allows them to assess whether the scenario is reasonable, which is another benefit of de-coupling the scenario selection from the choice of risk measures.

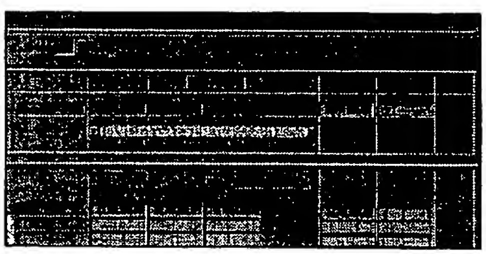

Figure 3.7: Stress Test Viewer

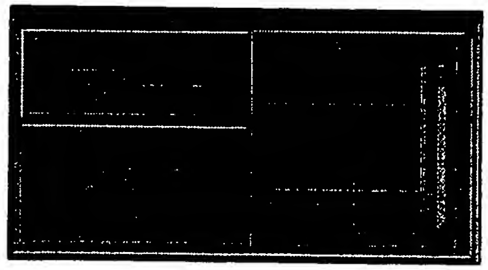

Figure 3.8: Scenario Viewer

Benchmarking the Mark-to-Future Implementation

As shown in Figure 3.9, HypoVereinsbank takes advantage of the additivity of a MtF approach to enhance performance in two ways. First, it can reduce elapsed time to compute risk measures by running multiple sessions of RiskWatch in parallel. Second, it can begin its overnight process before end-of-day positions are available. Setting aside the memory issue for the moment, the processing time required to run a single session would be prohibitive. If all goes well, it may take 10 hours: two hours of input/output (I/O) and eight hours of computation. Given that end-of-day positions are not available until 04:00, preliminary risk results are not available until the end of the next trading day (around 15:00). If corrections are required, final risk results can not be available until 02:00 the next morning. Without the MtF framework, the bank would need to completely rerun its overnight process to obtain the required risk measures—effectively doubling the processing time. This constraint puts daily risk measures at risk.

Given the size of HypoVereinsbank's portfolio (100,000 positions), the number of scenarios (4,000) and the number of risk factors (550), approximately 4.7 GB of memory are required for the positions and results alone, not including the memory required for the scenarios themselves. As such, given current technological limitations, running the computation in a single session is not feasible.

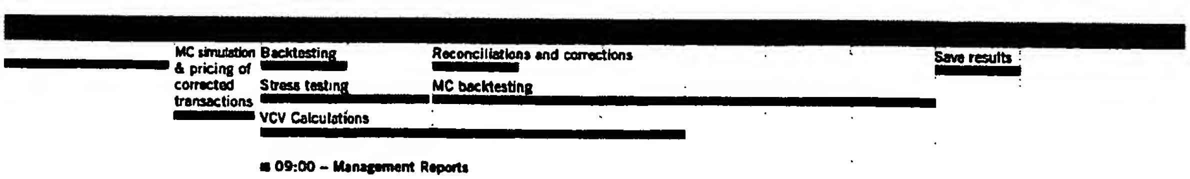

Table 2.2 summarizes the details of a comparison between two implementations, one a traditional, centralized approach outside the MtF framework and the other a centralized approach in the MtF framework.

Total Elapsed Compute Time (Base)	more than 10 hours	2-3 hours
Preliminary Risk Results Available	16:00	07:00
Final Risk Results Available	02:00 the next day	09:00
Memory Requirements	Not Feasible	Feasible

Table 2.2: Performance comparison of traditional and MtF implementations

Benefits of a Mark-to-Future System

As one of the first to implement the MtF paradigm, HypoVereinsbank has positioned itself well with respect to its industry counterparts in the field of risk measurement and management. Because MtF values can be computed once then used many times there is no need to re-compute MtF values in order to

- aggregate risk measures at different levels in a portfolio hierarchy and along multiple attributes
- compute statistics on the portfolio
- perform quick marginal VaR or *what-if* analyses.

This leads to

- improved performance in terms of elapsed time
- timely and actionable risk information
- effective limits management along multiple hierarchies and measures.

Special features of the HypoVereinsbank implementation enable risk managers to

- identify and view scenarios which cause selective risk results
- add user-defined scenarios easily.

Using the MtF framework, HypoVereinsbank has implemented a risk measurement system that enables it to achieve group-wide consistency of methodology and transparency of results as well as functional integration of various risk classes and P&L. The bank can use a number of different methods to measure risk and decompose risk at different levels in the portfolio. The necessary risk information is delivered in timely manner to those who must act upon the information.

Moving Forward

For HypoVereinsbank, the ultimate goal is to have the RiskWatch system linked to all of its investment banking products in its four major trading centers. The bank plans to have full functionality by the end of 2000 or early 2001; however, integration will be an ongoing process as new products and technology enhancements are developed. In the short term, the bank plans to extend its MtF implementation to include counterparty credit risk measurement and management. Further refinements to the MtF implementation are also under consideration.

As the bank proceeds with its credit risk implementation, it intends to adopt a longer time horizon and multiple time steps. It has already begun to investigate the appropriate time horizon and number and length of time steps for the credit risk measurement of its global portfolio. With longer time horizons, HypoVereinsbank must consider rebalancing using dynamic portfolio strategies and accounting for all events and their funding consequences. Using basis instrument mapping the bank can realize significant improvements in performance and storage.

HypoVereinsbank plans to move forward with its implementation of Algorithmics' MtF approach. From the bank's perspective, the MtF framework leapfrogs today's paradigm of risk capital as a static function of Value-at-Risk.®

Part 3 References

Bank for International Settlements, *Ammendment to the Capital Accord to Incorporate Market Risks*, Basle Committee on Banking Supervision, January 1996.
Available at http://www.bis.org/publ/

Longerstaey, J. and P. Zangari, 1996,
RiskMetrics-Technical Document, 4th ed., New York: Morgan Guaranty Trust Co. Available at
http://www.riskmetrics.com/rm/index.cgi

Penny, D., 1999,
'Technology trends,' *Algo Research Quarterly*, 2(2): 35-40.

Algorithmics Incorporated, 1999a,
Risk++ *User's Guide*, Version 3.1, Toronto.

Algorithmics Incorporated, 1999b,
RiskWatch *User's Guide*, Version 3.1, Toronto.

Part 4
Advanced Mark-to-Future Applications

An interesting selection of advanced applications that address important problems in risk management is provided. Topics include: the accurate measurement of counterparty risk, the integration of market and credit risk, the measurement of risk over long horizons using dynamic portfolios and the development of tools to optimally restructure portfolios. Each of these applications is built on a Mark-to-Future Cube.

4

This section presents a series of advanced applications of the Mark-to-Future framework. These applications demonstrate how this framework makes it computationally feasible to solve complex problems while integrating risks that would otherwise be treated separately and independently. The four selected papers consider important problems in risk management: the accurate measurement of counterparty risk, the integration of market and credit risk, the measurement of risk over long horizons and dynamic portfolios and the development of tools to optimally restructure portfolios.

The first two papers present credit risk measurement applications. There are two types of credit risk measurement models: counterparty credit exposure models and portfolio credit risk models. Derivative desks traditionally manage credit risk by monitoring and placing limits on counterparty credit exposures. Since their main focus is on risk at the counterparty level, counterparty credit risk models do not attempt to capture portfolio effects such as the correlation between counterparty defaults. In contrast, portfolio credit risk models measure credit capital and are specifically designed to capture portfolio effects and correlations of obligor credit quality changes.

The first paper, by Aziz and Charupat, focuses on the estimation of counterparty credit risk for derivative portfolios. The authors introduce the measurement of counterparty credit risk in a MtF framework based on multi-step Monte Carlo simulation. The MtF exposure profiles are contrasted to the capital adequacy requirements stipulated by the Basle Accord. By simulating counterparty portfolios through time over a wide range of scenarios, the MtF methodology accurately captures the term structure of credit risk and explicitly accounts for the contingency of derivative portfolios and credit risk on the market. MtF can accurately model natural offsets, netting, collateral and various mitigation techniques used in practice. Furthermore, the MtF method can explicitly incorporate probabilities of default and recovery in an efficient computational manner to obtain counterparty credit losses. Finally, it is shown that the MtF infrastructure provides a much richer set of information with which credit managers can characterize exposures and losses as well as to explain their causes.

The second paper, by Iscoe, Kreinin and Rosen, presents an integrated model of market and credit risk. This is a novel multi-step portfolio credit risk model that integrates counterparty exposure simulation and advanced portfolio credit risk methods. The integrated model overcomes a major limitation currently shared by portfolio models, namely accounting for the credit risk of portfolios when exposures depend on market levels. This is the case when a portfolio contains derivatives such as swaps, forwards, options or loans with embedded optionality. The model is computationally efficient because it combines a MtF framework of counterparty exposures and a conditional default probability framework. This approach has a number of computational benefits. First, the number of scenarios requiring expensive portfolio valuations is minimized. Second, the model is based on the same computations used for

monitoring counterparty exposures. A MtF framework allows users to exploit these computations and use them for both counterparty exposures and portfolio credit risk. This also leads to more consistent enterprise risk measurement. Finally, advanced Monte Carlo or analytical techniques that take advantage of the problem structure can be used to solve the problem more efficiently than standard Monte Carlo methods.

The third paper, by Black and Telmer from the Bank of Canada, describes a practical application of dynamic portfolio strategies in a MtF framework. The authors show the necessity of such a framework to analyze the problem of debt issuance in a stochastic environment. Standard parametric VaR methodologies are not appropriate for these problems, where risk must be measured over long time horizons and portfolio effects such as issuance, aging, coupon payments, maturities and settlement must be taken into account. In contrast, a MtF framework with dynamic portfolio strategies naturally determines how the interest expenses of a liability portfolio evolve over time and permits calculation of the cost versus risk trade-off of various strategies. These strategies may vary from simple, fixed rules to dynamic rules that are contingent on future events.

The last paper, by Mausser and Rosen (1999), develops risk management tools in a MtF environment. While financing institutions have devoted significant effort to measuring risk, the development of tools to understand and manage risk has generally lagged behind. Risk management is a dynamic endeavour, necessitating tools that identify and reduce the sources of risk, and ultimately optimize the utilization of available financial products to obtain the desired risk profiles. MtF risk management and optimization tools offer many advantages over standard approaches, such as those based on a mean-variance framework. This is especially true when the portfolio contains instruments with non-linear pay-offs, when the market distributions are not normal, or when there are multiple horizons. In particular, MtF tools are very useful not only for market risk, but also for credit risk, where the exposure and loss distributions are generally skewed and far from normal.

Exposure Loss

We report a study of the estimation of credit exposure and credit loss of a portfolio of derivative transactions. The estimation is performed using a Monte Carlo simulation. The results are compared to the exposure and capital reserves obtained under the method recommended by the Bank for International Settlements (BIS). We show that the simulation method provides a much richer set of information for credit risk managers. Also, depending on the current exposure and the nature of the transactions, the BIS method can fail to account for potential exposure. In addition, depending on the level or prudence one requires in setting capital reserves, the BIS reserves can be either excessive or grossly insufficient.

The dramatic increase in over-the-counter derivative trading during the past ten years has made credit risk a major concern for banks. With increased competition and tighter spreads, banks must accurately quantify the credit risk they are facing so that they can appropriately price their products, set the proper level of capital reserves and manage their credit line efficiently. Currently, different banks employ different credit risk measurement techniques, ranging from applying a fixed percentage to the notionals of the transactions, to creating distributions of future credit exposure and credit loss.

This paper focuses on estimating counterparty credit risk for derivative portfolios. We accomplish three objectives: introduce the Monte Carlo Simulation method used to measure counterparty credit risk; highlight several significant aspects of counterparty exposure and loss profiles that result from Monte Carlo simulation using an example; and contrast these results to the Credit Equivalent Amounts and capital adequacy requirements stipulated by the Basle Accord (BIS 1988).

In the first section we define several credit risk terms used throughout this paper. Next, we outline the BIS and Monte Carlo-based methods used to estimate counterparty credit exposures and losses. We then compare the two methods, underscoring the limitations of the factor-based approach to measure capital and manage portfolios. The results are outlined followed by the conclusions we draw from using these two approaches.

Credit Risk Definitions

In this paper, we define **credit risk** as the risk of loss that will be incurred in the event of default by a counterparty.[1] Default occurs if the counterparty fails to honour its contractual payments.

Calculating potential credit losses requires the modeling of three processes that define: (i) the counterparty credit exposure at the time of default; (ii) the probability that default will occur; and (iii) the amount that can be recovered after default occurs.

There are two types of credit risk: pre-settlement risk and settlement risk. Pre-settlement risk refers to losses arising from prematurely terminating a contract due to counterparty default. This loss depends on the replacement value of the contract at the time of default. In contrast, settlement risk refers to the losses realized when the contractual payments are not received on the settlement date. In this paper, we consider only pre-settlement risk.

Credit Exposure

Credit exposure is the cost of replacing or hedging the contract at the time of default. This is the maximum value that will be lost if the counterparty to that contract defaults. Since default is an uncertain event that can occur at any time during the life of the contract, we consider not only the contract's current credit exposure, but also potential changes in the exposure during the contract's life. This is particularly important for derivative contracts whose values can change substantially over time and according to the state of the market. For this reason, we introduce three measures of credit exposure: Actual Exposure, Total Exposure and Potential Exposure.

We define the **Actual Exposure** of contract c at time t as the maximum of zero and the value of the contract at that time:

$$AE(c, t) = max\{0, V(c, t)\} \tag{1}$$

where $V(c, t)$ is the value of contract c at time t. As defined, $AE(c, t)$ is the maximum amount that will be lost if default occurs and if the contract is replaced at time t. If t is the current time, the Actual Exposure then depends on the current marked-to-market value of the transaction. If t is a future point in time, then $AE(c, t)$ is defined with respect to an assumed state of the market at time t.

The **Potential Exposure** of the contract at time t is the maximum *additional* amount (over the Actual Exposure) that will be lost if default occurs not at time t, but at some time τ, between t and the maturity of the contract, T:

$$PE(c, t) = max\{0, max_{t < \tau \le T}\{PV_t[V(c, \tau)] - V(c, t)\}\} \tag{2}$$

where $PV_t[\cdot]$ is a function that transforms future values to their present values at time t. A **market scenario** is the path that the market takes during the period from t to T. The Potential Exposure must be defined with respect to an assumed market scenario. This is because the value of the contract at any point in time during that period depends on the values that the underlying market risk factors take. As such, the Potential Exposure takes into account the aging of the portfolio and possible adverse movements in the underlying variables under that market scenario.[2]

The **Total Exposure** of the contract at time t is the sum of the Actual Exposure and the Potential Exposure at that time:

$$TE(c, t) = AE(c, t) + PE(c, t) \tag{3}$$

Therefore, it is the maximum value that will be lost on the contract if the counterparty defaults at any point from time t to the end of the contract's life, under an assumed market scenario.

As an example, consider a pay-fix/receive-float interest-rate swap whose current marked-to-market value is negative. In this case, the contract's current Actual Exposure is zero. However, if interest rates rise in the

future, the contract can have a positive value. Although the contract has no current credit exposure, it carries potential additional exposure that presents credit risk to the holder. The Total Exposure captures this future positive value.

At the counterparty level, credit exposure depends on the netting arrangement and other credit mitigation provisions. For example, if full netting is allowed, contracts with positive values can be offset by contracts with negative values to reduce the net exposure. In this case, the Actual Exposure to counterparty P at time t is the maximum of zero and the sum of the values of all contracts with that counterparty:

$$(4) \quad AE(P,t) = max\left\{0, \sum_{c \in P} V(c,t)\right\} = max\{0, V(P,t)\}$$

The Potential Exposure to a counterparty under an assumed market scenario is the maximum additional exposure at some future time. It can be obtained by substituting the portfolio's value, $V(P,t)$, for the contract's value, $V(c,t)$, in Equation 2. The Total Exposure is the sum of the counterparty's Actual Exposure and its Potential Exposure.

On the other hand, if netting is not allowed, the Actual Exposure to a counterparty is simply the sum of the values of all positive-valued contracts with that counterparty. In other words, it is the sum of each contract's Actual Exposure:

$$(5) \quad AE(P,t) = \sum_{c \in P} max\{0, V(c,t)\} = \sum_{c \in P} AE(c,t)$$

The Potential Exposure in this case is conservatively defined as the sum of each contract's Potential Exposure at that time:

$$(6) \quad PE(P,t) = \sum_{c \in P} PE(c,t)$$

Note from Equation 2 that the Potential Exposure of a contract depends on the maximum value the contract will take during its life, and that different contracts may reach their maximum values at different points in time. Therefore, by defining the Potential Exposure to the counterparty in this manner, we allow for the conservative possibility that in the absence of a netting agreement, the counterparty may cherry-pick the timing of default of each contract.

Finally, from Equations 3, 5 and 6, the Total Exposure to the counterparty in the absence of a netting agreement is simply the sum of each contract's Total Exposure:

$$(7) \quad TE(P,t) = \sum_{c \in P} AE(c,t) + \sum_{c \in P} PE(c,t) = \sum_{c \in P} TE(c,t)$$

Credit Losses

While credit exposure is the maximum amount that will be lost if the counterparty defaults, the credit losses take into account the amount that can be recovered after the default occurs. That is, credit loss on contract c, if default occurs at time τ under some assumed market scenario is:

$$L(c,\tau) = AE(c,\tau) \times [1 - R(c,\tau)]$$

where the recovery rate, $R(c,\tau)$, is the percentage of the value that will be recovered on the contract if default occurs at time τ. Usually, $R(c,\tau)$ depends on the seniority of the contract. $R(c,\tau)$ can also depend on the state of the market when default occurs.

If netting is allowed, counterparty credit losses are a function of Actual Exposure, which in turn depends on the netting provisions stipulated in the master agreement. The credit loss on a counterparty portfolio if default occurs at time τ is:

$$(8) \quad L(P,\tau) = AE(P,\tau) \times [1 - R(P,\tau)]$$

When netting is allowed, a single recovery rate is applied to all transactions. This is because it is reasonable to assume that the contracts for which netting is allowed have the same seniority class.

If, on the other hand, netting is not allowed, credit losses from a counterparty portfolio (if default occurs at time τ) are conservatively defined as a function of the portfolio's Total Exposure:

$$(9) \quad L(P,\tau) = \sum_{c \in P} L(c,\tau) = \sum_{c \in P} TE(c,\tau)[1 - R(c,\tau)]$$

The Total Exposure is used because, in the absence of a netting agreement, the institution cannot offset the exposure immediately after default occurs on the first contract. Equation 9 allows for the possibility that the counterparty may choose to default on the contracts when each of them reaches its maximum exposure. In this case, it is possible to have contracts with different recovery rates.

In practice, master agreements can be quite complex. For example, each counterparty may hold several different master agreements that permit netting across certain instruments, but not others. In addition, netting may not be applicable in certain jurisdictions at the time of default. An accurate estimation of counterparty exposures requires that netting agreements be accurately modeled in a flexible netting hierarchy. Furthermore, given the strong dependence of exposures to the netting provisions, stress testing these provisions is very important.

Netting is perhaps the most popular mitigation technique. Other mitigation techniques include the posting of collateral, marked-to-market caps, re-couponing and early termination clauses (Wakeman 1997). Accurate exposure calculations require a thorough modeling of these provisions.

Credit Risk Measurement Methodologies

Current credit risk measurement methodologies differ from one another in their assumptions of (i) credit exposure, (ii) default probability, and (iii) recovery rate. A simple methodology may assume that the contract's credit exposure is equal to its notional. Potential loss is then estimated by applying a fixed percentage to the notional, where, presumably, this percentage encompasses both the default possibility and the recovery rate. A more sophisticated methodology creates distributions of future exposures, default rates and recoveries before calculating the losses through simulation.

In this section, we introduce two methodologies. The first is the method put forth by the Bank for International Settlements. It is currently the only method financial institutions are allowed to use to calculate their capital reserves. The second method is a simulation-based method in which Monte Carlo simulations are employed to create profiles of exposures and losses.

The BIS Methodology

Concerns about credit risk have prompted the Bank for International Settlements to introduce capital adequacy requirements for financial institutions that deal in derivative securities. This requirement, commonly referred to as the Basle Accord (BIS 1988), specifies the method that banks must use to calculate credit exposure and is intended to establish a minimum level of capital reserves. Though the methodology prescribed by the BIS follows the framework described in the previous section, some characteristics of the methodology are noteworthy.

First, the Actual Exposure, and thus the Potential and the Total Exposure, are defined only at the current time, $t=0$. At the contract level, the Total Exposure of a derivative position, known as the Credit Equivalent Amount (CEA), consists of two parts, the Actual Exposure and the Potential Exposure. The Actual Exposure is defined according to Equation 1. The Potential Exposure, on the other hand, is calculated by multiplying the notional of the transaction, N, by a pre-defined credit conversion factor, CF:

$$PE = CF \times N \tag{10}$$

This factor depends on the broadly classified types of the securities underlying the derivative contracts and the time to maturity of the contracts. Table 1 summarizes the factor by the maturity date and the type of underlying.

The credit conversion factor is intended to account for the possibility that future exposures may exceed current exposures. The Potential Exposure, calculated according to Equation 10, is known as the 'static add-on for Potential (Future) Exposure' because it is a fixed amount of the notional, independent of passing time, that is added to the Actual Exposure. Thus, the CEA:

$$CEA = AE(c, 0) + PE(c, 0) \tag{11}$$

is defined only at the current time.

If netting is allowed, the CEA is equal to the sum of the netted Actual Exposure (Equation 4) and the sum of each transaction's Potential Exposure (Equation 10) adjusted by a netting factor, NF:

$$CEA = AE(P, 0) + NF \times PE(P, 0) \tag{12}$$

The netting factor, NF, is defined as:

$$NF = (0.4 + 0.6 * NGR)$$

where the Net-to-Gross Ratio, NGR, is the ratio of the Actual Exposure with netting (Equation 4) to the Actual Exposure without netting (Equation 5). Note that if netting is not permitted, $NGR=1$, $NF=1$ and the CEA for the portfolio is the sum of each transaction's CEA.

< 1 year	0%	1%	6%	7%	10%
1-5 years	0.5%	5%	8%	7%	12%
> 5 years	1.5%	7.5%	10%	8%	15%

Table 1: BIS credit conversion factors, CF (BIS 1997)

The Capital Reserves, CR, are then calculated by multiplying the CEA first by a risk-weight factor, RF, which depends on the type of the counterparty to the contract, and then multiplying by an 8% capital-to-exposure ratio:

$$CR = RF \times CEA \times 0.08 \tag{13}$$

The BIS guideline for the values of the risk-weight factors are summarized in Table 2.

Organization for Economic Cooperation and Development (OECD) governments	0%
OECD banks and public-sector entities	20%
Corporate and other counterparties	50%

Table 2: Risk weights for off-balance-sheet transactions, *RF* (BIS 1997)

The foremost conclusion of the ISDA evaluation of credit risk and regulatory capital requirements (ISDA 1998) is that there is an urgent need for reform. The report promotes adoption of a models-based approach as an alternative to the current standardized rules. A models-based approach, such as the MC approach described below, is more consistent with internal risk management practice and is more conducive to prudent credit risk management. They note a number of shortcomings in the current regime, including

- limited differentiation of credit risk among broad categories of credit
- static measures of default risk based on an 8% capital requirement adjusted for broad categories of risk
- no recognition of the term structure of credit risk, and thus no recognition of the evolution of risk factors and the greater probability of default associated with longer exposures
- simplified potential future risk calculation, leading to limited and inexact recognition of netting and of moneyness of the position
- lack of recognition of portfolio diversification effects and thus encouragement for judicious diversification.

While the BIS assumes a static add-on for Potential Exposure, the simulation-based method computes the contract's value through time under each scenario. Moreover, the probability of default and recovery rate are implicitly taken into account by the BIS through the risk-weight factors and the capital-to-exposure ratio, whereas they are explicit under the simulation-based approach. In the next section, we consider the simulation-based method.

The Monte Carlo-based Method

Monte Carlo simulation is a comprehensive method for estimating credit exposures and losses for derivative portfolios. Although Monte Carlo simulation is generally difficult to implement and computationally intensive, it can realistically incorporate the impact of all sources of risk, offsetting correlations between various positions and counterparties, as well as netting and mitigation provisions.

In its most general form, during a Monte Carlo simulation a large number of joint scenarios are generated based on (i) the market risk factors affecting the value of the portfolio, (ii) credit events such as default and credit migration, and (iii) recovery rates for each contract upon default. Every scenario is a path over time of tens, and possibly hundreds, of risk factors covering market and credit events. Typically, the simulation is performed over the entire life of the transactions in the portfolio. Under each scenario and time step, the portfolio is revalued and recoveries are applied when default occurs. Finally, the results of the simulation are aggregated and various statistics are computed from the distributions of the exposures over time and from the distribution of losses.

A naive simulation is not computationally practical. A prohibitively large number of scenarios is required since credit events, especially default, occur with low probabilities.

This limitation can be overcome by exploiting some basic properties of the problem and by making some simplifying assumptions. In particular, one can speed up the computation dramatically by applying a simple 'decomposition' of the problem. We have defined credit exposure as the cost of replacing or hedging a contract when default occurs. Therefore, exposures depend only on the state of the market at that time, and not on the credit state of the counterparty.

As a first stage, one can run a simulation on the whole portfolio to compute a 'table of counterparty exposures'. The table of exposures summarizes the Actual Exposure (Equation 4 or 5) and Total Exposure (Equation 3 or 7) of every counterparty, under each market scenario and time step. This is generally the most computationally expensive step in the simulation since it requires the valuation of all the instruments in the portfolio under each scenario. Then, to compute credit losses one runs a new simulation with joint random draws from the table of exposures, counterparty defaults and recoveries. A large number of scenarios can be drawn in this case, since every time a default occurs in a scenario, the losses are simply computed by reading the table and applying the recovery rates; no expensive portfolio valuations are necessary.

The simulation of credit events can be eliminated altogether when computing counterparty credit losses if we further assume that (i) the probabilities of default and recovery rates are deterministic and (ii) market and credit events are independent.

The first assumption is quite common in many pricing models (Jarrow and Turnbull 1995, Hull and White 1995, and Jarrow et al. 1997). Das and Tuffano's model (1997) relaxes the assumption of deterministic recovery rates in Jarrow et al. and correlates recoveries to interest rates. Credit risk models such as Wilson (1997) and CreditMetrics (JPMorgan 1997) have also traditionally assumed deterministic defaults and recoveries, although this assumption can be relaxed. The CreditRisk+ model (Credit Suisse 1997) explicitly incorporates stochastic probabilities of default.

The assumption of market and credit independence, while not always realistic, greatly simplifies the calculations. Hence, it has been frequently employed; for example, Hull and White (1995) and Jarrow and Turnbull (1995) use it to price derivative securities, while Jamshidian and Shu (1997) use it in their credit risk calculations. Hull and White argue that the assumption is not unreasonable when the counterparties are large, well-diversified financial institutions whose books are unlikely to be sensitive to movements in a single market factor. The assumption is less realistic, however, when the counterparties are production firms that engage in derivative transactions to hedge their positions or to speculate in the areas in which they have special knowledge. Duffee (1996) demonstrates by example that in these cases, the assumption significantly impacts the results.

If one makes assumptions (i) and (ii) above, the loss distribution can be obtained explicitly, without further simulation, from the table of exposures, the recoveries and the cumulative default probabilities of the counterparty. This is depicted in Figure 1 for a single counterparty.

Consider a set of n market scenarios, ω_j, each with probability p_j, and assume that default can occur at m discrete points in time, t_i. In this figure, $\lambda(i,j)$ is the conditional loss if default occurs at time t_i and scenario ω_j, i.e. $\lambda(i,j)=L(P,t_i)$, as given by Equation 8 or 9, assuming scenario ω_j occurred. These are the entries of the table of exposures adjusted by the recovery rates. Because of the independence assumption, the probability of a loss of $\lambda(i,j)$ is simply $p_j \times D_i$, the product of the probability of market scenario ω_j occurring and the probability of default at time t_i. The full counterparty loss distribution is simply obtained by ordering these losses in ascending order.

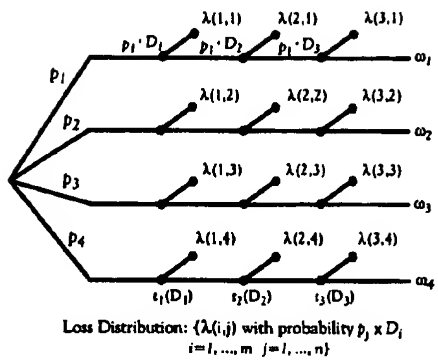

Figure 1: Calculation of losses in MC simulation

Figure 1 illustrates a small example with three points in time and four market scenarios, resulting in 13 possible loss values. The last value corresponds to the case of no default (and thus zero loss) occurring before the last time point.

Based on these loss numbers and their associated probabilities, various statistics can be obtained. We define three measures of credit loss: the Expected Loss, and two measures defined as the maximum losses that can be expected to occur at some level of confidence, α.

- Expected Loss, *EL*

  The Expected Loss from a given time, t_k, (across scenarios and time):

  $$(14)\quad E_{D,\omega|\tau>t_k}[L(P,\tau)] = \sum_{j=1}^{n} \sum_{i=k+1}^{m} \lambda(i,j) \times p_j \times D_i$$

  where $E_{D,\omega|\tau>t_k}[\bullet]$ is the expectation function with respect to both the scenario and default probabilities, conditional on default not occurring before t_k.

- Maximum Scenario Loss, $MSL(\alpha)$

  The **scenario credit loss** is the expected loss from a given time, t_k, to the contract's maturity, T, under an assumed market scenario j:

  $$(15)\quad E_{D|\tau>t_k,\omega_j}[L(P,\tau)] = \sum_{i=k+1}^{m} \lambda(i,j) \times D_i$$

  where $E_{D|\tau>t_k,\omega_j}[\bullet]$ is the expectation function, in the interval (t_k, T), with respect to the default probability, conditional upon market scenario ω_j, and default not occurring before time t_k.

The Maximum Scenario Loss at a level of confidence level α, $MSL(\alpha)$, is obtained by ranking the expected value of the losses under each scenario (Equation 15) and choosing the maximum (at level α):

$$\text{(16)} \quad Pr\{E_{D|\tau > t_i, \omega_j}[L(P,\tau)] \leq MSL(\alpha)\} = \alpha$$

- Maximum Loss, $ML(\alpha)$

  The Maximum Loss, across time and scenarios with respect to a confidence level, α, $ML(\alpha)$, can be obtained directly from the distribution of the losses at all possible combinations of time points and market scenarios, ranking them and choosing the one that corresponds to the value at which the chance of losses higher than $ML(\alpha)$ does not exeed $1-\alpha$:

$$\text{(17)} \quad Pr\{\lambda(i,j) \leq ML(\alpha)\} = \alpha$$

In the case study in the next section, we assume that the probabilities of default and recovery are deterministic, and that counterparty defaults are independent of the market risk factors. The table of counterparty exposures is computed from a multi-step Monte Carlo simulation, where scenarios for the market risk factors (interest and foreign exchange rates) follow a general multi-curve, multi-factor model with mean reversion.

Mean reversion prevents the volatilities of the underlying processes from becoming unduly large even though a credit risk analysis covers a long simulation time horizon.

Each node on an interest rate curve is treated as a separate risk factor. The scenarios are then generated based on a historical estimation of the volatilities and correlations of the risk factors. These volatilities and correlations are assumed to be constant over time.

A Case Study

In the previous section we introduce the MC method. In this section we accomplish the remaining two objectives of this paper. First, we present the details of a test portfolio for which we produce a credit risk report based on a Monte Carlo credit analysis. We examine this report for the information that might be derived by a credit risk manager and highlight several significant aspects of the counterparty exposure and loss profiles. Finally, we contrast these results with those derived from the BIS approach.

The Sample Portfolio

The sample portfolio of the Spadina Bank is constructed such that it is representative of derivative portfolios held by major North American banks. It contains portfolios for six counterparties with four types of derivative transactions—foreign exchange options, foreign exchange forwards, currency swaps and interest-rate swaps. Five of the counterparties are OECD-incorporated banks with credit ratings of A or higher. The other counterparty is a BB-rated corporation. The counterparties, together with their credit ratings and outstanding contracts as of June 4, 1997, are summarized in Table 3.

The marked-to-market values of the transactions are calculated using the data on interest rates and foreign exchange rates as of June 4, 1997. These data are obtained from Reuters on-line services. In total, there are sixteen currencies (including USD) in the portfolio.

The Spadina Bank has a netting agreement with every counterparty. All agreements allow cross-product netting, except for the agreement with Ruby which allows all the foreign exchange contracts (i.e., FX forwards, FX options and currency swaps) to be netted together, and all interest-rate contracts (IR swaps) to be netted together, but does not permit netting between the two types.

Monte Carlo Simulation

To generate the distributions of credit exposures and credit losses at future points in time, we create 1,000 scenarios on the 16 relevant interest rate (zero) curves and 15 foreign exchange rates using a multi-factor, multi-step Monte Carlo method with mean reversion.

Because the volatilities of some Asian interest rates and currencies are extremely high, they are adjusted such that none of them exceeds 8%. This is to ensure that the simulation does not generate values of risk factors that are unrealistically high. We believe that this adjustment is justified based on past behaviors of interest rates. For interest rates that have low volatilities such as those of the G7 countries, a mean reversion coefficient of 0.01 is used. For other interest rates, we use a coefficient of 0.1.

Probabilities of Default and Recovery Rates

To calculate credit losses, we require the probabilities of default, D_i, and the recovery rates, $R(c,t)$, should default occur. The default probabilities are calculated based on average one-year transition matrices for banks and for corporations (Standard & Poor's 1997). These transition matrices contain the historically average probabilities that a company (bank or corporation) with a certain credit rating migrates to another rating or default by the end of one year. The cumulative probability that a company will default by the end of n years is obtained by multiplying the appropriate transition matrix by itself n times. These cumulative probabilities are summarized in Table 4.

The probability of default in a future period can be determined from Table 4. For example, the probability that a AA bank will default between year 3 and year 5 *and* not default before year 3 is 0.06% (= 0.08% – 0.02%).

Note that our calculation of default probabilities assumes that the default processes are Markovian. That is, we assume that the probability of migration depends solely on the counterparty's current credit rating, and not on its credit rating history. A conservative recovery rate of zero is assumed for all contracts.

Diamond (Bank, A)						
FX Forward	525	-12.177	5%	1.8	4.8	5,482.403
FX Options	298	6.348	-3%	1.7	5.0	23.750
IR Swaps	286	-233.334	98%	5.7	19.1	11,990.292
TOTAL	1,109	-239.163		2.1	19.1	17,496.445
Ruby (Bank, AAA)						
FX Forward	16	-0.271	0%	0.9	3.4	165.679
Currency Swaps	46	432.047	83%	0.4	15.4	4.106.855
IR Swaps	93	91.310	17%	5.3	15.9	3,404.842
TOTAL	155	523.086		4.7	15.9	7,677.365
Sapphire (Bank, AA)						
FX Forward	535	44.193	-70%	1.8	4.8	5,594.322
IR Swaps	594	-89.991	170%	4.9	16.0	17,897.094
TOTAL	1,146	-45.799		3.2	16.0	23,491.416
Topaz (Bank, A)						
FX Forward	12	-8.534	-6%	1.4	3.9	133.581
IR Swaps	21	147.396	106%	5.1	9.9	3,042.508
TOTAL	33	138.862		4.0	9.9	3,176.089
Emerald (Corp, BB)						
FX Forward	117	-10.272	100%	1.4	3.9	1,248.488
TOTAL	117	-10.272		1.4	3.9	1,248.488
Turquoise (Bank, AA)						
FX Forward	455	-65.495	121%	2.0	3.9	4,710.904
IR Swaps	26	9.062	-21%	6.1	15.9	956.849
TOTAL	481	-56.433		2.1	15.9	5,667.753
Total Portfolio	3,024	310.281		2.4	19.1	58,757.566

Table 3: Sample portfolio by counterparties. The notionals of the swaps are equal to the swap principals. The notionals of the foreign exchange forwards and foreign exchange options are obtained by multiplying the contract size by one unit of the underlying currency, and then converting it into USD using the exchange rates on June 4, 1997.

AAA—Bank	0.00	0.00	0.02	0.09	0.36	0.91	1.81
AA—Bank	0.00	0.02	0.08	0.33	1.10	2.43	4.25
A—Bank	0.04	0.17	0.48	1.38	3.31	5.80	8.60
BB—Corporation	1.10	4.62	9.02	15.88	24.22	31.19	36.92

Table 4: Cumulative probabilities of default

Results

Based on the implementation of a MC simulation described above, we produce a credit risk report which we examine for information that might be derived by a credit risk manager. We highlight several significant aspects of the counterparty exposure and loss profiles and contrast these results with those derived from the BIS approach.

MC-based Credit Exposures and Losses

Figure 2 is an example of a credit risk report for the Spadina Bank that summarizes the results of an analysis based on the MC approach. It comprises two tables. The first contains statistics on credit exposure:

- Actual Exposure, *AE*, (Equation 1)
- Expected Total Exposure, *ETE*
- Maximum Total Exposure at a 99% confidence level, *MTE(99%)*.

The Total Exposure to a counterparty under a scenario is calculated using Equation 3 or 7, depending on the provisions for netting.[3] The Maximum Total Exposure, *MTE(99%)*, is then obtained by ranking the Total Exposures under all 1,000 scenarios and choosing the one for which there is a less than 1% chance that Total Exposure in other scenarios will be higher than MTE(99%):

$$Pr\{TE(P,t) \leq MTE(99\%)\} = 99\%$$

The second table summarizes the results of the loss calculations:

- Expected Loss, *EL*, (Equation 14)
- Maximum Scenario Loss at a 99% confidence level, *MSL(99%)*, (Equation 16)
- Maximum Loss at a 99% confidence level, *ML(99%)*, (Equation 17)
- Maximum Loss at a 99.9% confidence level, *ML(99.9%)*, (Equation 17).

As well, the report includes a set of graphs that illustrates the evolution of exposures over time for each counterparty. Each of the figures contains three different profiles of Actual Exposure through time:

- Expected Actual Exposure[4]
- Maximum Actual Exposure at a 95% confidence level, *MAE(95%)*
- Maximum Actual Exposure at a 99% confidence level, *MAE(99%)*.

The report is dated June 4, 1997. We review the information that a credit risk manager might extract from this summary credit report, beginning with an examination of the credit exposures.

Only two counterparties—Ruby and Topaz—have positive current Actual Exposures. The Total *AE* represents the portfolio's replacement cost should both counterparties default today. The portfolio's current value of 310 million USD is about half of its current Actual

The Spadina Bank Credit Report June 4, 1997

Diamond (A)	-239.16	0.00	24.61	183.46
Ruby (AAA)	523.08	523.08	791.05	1,233.93
Sapphire (AA)	-45.79	0.00	22.07	92.70
Topaz (A)	138.86	138.86	185.44	298.54
Emerald (BB)	-10.27	0.00	0.41	10.14
Turquoise (AA)	-56.43	0.00	23.06	75.63
TOTAL	310.28	661.94	1,046.56	1,894.42

Diamond (A)	0.18	1.80	1.88	38.61
Ruby (AAA)	0.34	0.70	0.00	76.68
Sapphire (AA)	0.08	0.48	0.00	25.62
Topaz (A)	0.49	1.00	0.00	134.31
Emerald (BB)	0.01	0.04	0.00	0.00
Turquoise (AA)	0.05	0.29	0.00	14.99

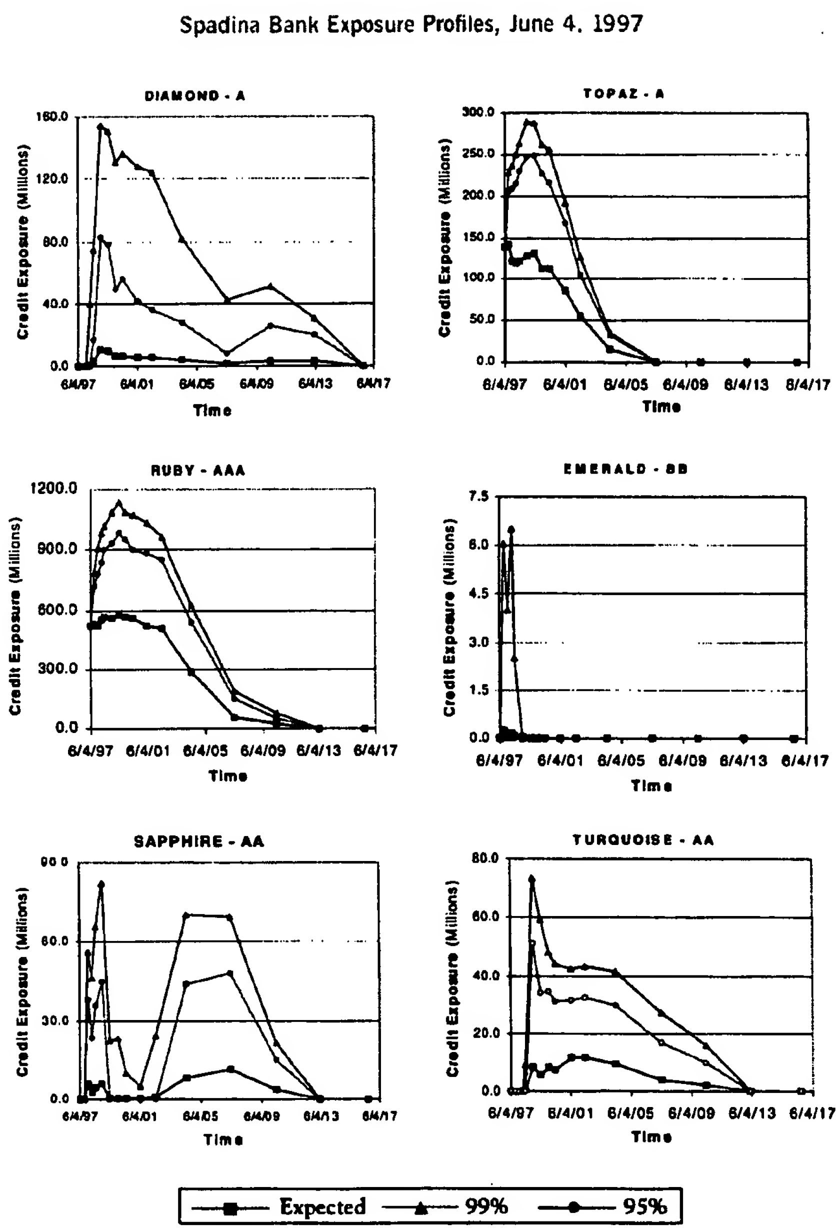

Figure 2: Spadina Bank credit report and exposure profiles

Exposure (662 million USD) because netting across counterparties is not permitted.

The tabulated Expected Total Exposure and Maximum Total Exposure by counterparty suggest that future exposures are greater than current exposure in each case. This suggestion is confirmed by the counterparty exposure profiles. Each of the three exposure profiles for each counterparty increases from the current value and remains positive for some period during the simulation. At the portfolio level, Expected and Maximum Total Exposures, 1,047 million USD and 1,894 million USD respectively, suggest that future credit exposure is, on average, 58% higher than the current exposure. The Maximum Scenario Loss, *MSL(99%)*, is almost three times higher.

The exposure profiles identify the periods in which counterparty default would be most financially damaging. Consider the exposure profiles of Diamond. The Maximum Actual Exposures, *MAE(95%)* and *MAE(99%)*, are zero during the first six months, and then surge to peaks of about 82 million USD and 150 million USD, respectively. Using the simulation infrastructure to investigate the causes of this increase, the credit manager discovers that it is primarily the result of foreign-exchange forward and option contracts, whose values drastically increase in a number of scenarios due to the volatilities of their underlying currencies. The behaviour of the Expected Actual Exposure, however, suggests that there are other scenarios where the values of these contracts decrease; otherwise, the mean exposure would sharply increase as well.

Next we consider the report on credit losses, beginning with a comparison of the maximum credit losses at 99% and 99.9% confidence levels. The *ML(99%)* is non-zero only for Diamond. Default is such a low-probability event that (i) for *each* of the five counterparties, there is a less than 1% chance that it will default *when* the exposure to it is positive, and (ii) there is a less than 1% chance that Diamond will default when its exposure is greater than 1.88 million USD.

At an increased level of confidence of 99.9%, *ML(99.9%)*, we note that all counterparties except Emerald have positive losses, and that the maximum credit loss for Diamond is no longer the largest. Let's analyze the losses of these two counterparties.

Although Emerald is a BB-rated corporation, its *ML(99.9%)* is zero for two reasons. First, all of its contracts expire within four years, half of them in two years. Secondly, its contracts are collectively out-of-the-money (i.e., the current aggregated marked-to-market value is negative). As a result, the simulated exposure profiles over the lives of its contracts are zero at most time points and under most scenarios. Thus, its *MSL(99%)* and its Expected Loss are low.

The Maximum Scenario Loss for Diamond is the highest even though its exposures are much lower than those of Ruby and Topaz. Compared to Ruby, which is a AAA bank, Diamond has a higher probability of default because it has a lower credit rating. This causes Diamond's *MSL(99%)* to be higher. Compared to Topaz's, Diamond's contracts have much longer maturities. As a result, the calculation of Diamond's *MSL(99%)* (Equation 16) spans a much longer period. The longer period is especially important because, from Table 4, the marginal probabilities of default are higher in later years. Therefore, even though Diamond and Topaz have the same credit rating, there are a few scenarios under which Diamond has a higher Maximum Scenario Loss as the effect of Diamond's longer maturities dominates the effect of Topaz's higher credit exposures in the early years.

Although Diamond has the highest *MSL(99%)*, it does not have the highest Expected Loss. This is because its Expected Actual Exposure is only slightly positive through time. The Expected Losses of Ruby and Topaz are higher, which is to be expected since their exposure profiles indicate prolonged, significant, positive Expected Actual Exposures.

Though the exposures summarized in the report indicate that substantial losses can be incurred, the probabilities that those losses will occur, and thus the Expected Losses, are substantially smaller.

Comparison of Methods

In this section, we compare measures of credit exposure and loss under the MC approach to the Credit Equivalent Amounts, CEAs, and Capital Reserves, CR, under the BIS methodology.The issues raised by ISDA and noted above are well-illustrated by this case study.

The ratios of Maximum Total Exposures, *MTE(99%)*, to CEAs are depicted in Figure 3. They range from 0.56 for Emerald to 1.73 for Topaz, suggesting that the CEAs can be higher or lower than the simulated *MTE(99%)*, depending on the nature of the contracts with each counterparty. In general, it appears that the CEAs are higher when the counterparty's current Actual Exposures are zero (i.e., Sapphire, Emerald and Turquoise, but not Diamond). This is to be expected since the BIS approach applies the same add-on Potential Exposure regardless of the moneyness of the current position. As a result, unless the exposure increases substantially over time (which is the case with Diamond), the CEAs are likely to be higher than the simulated Maximum Total Exposure.

Because losses are a function of exposures, the discrepancy between the two methods is propagated through the loss calculation. Figure 4 shows the log ratios of four different MC measures of credit losses to the BIS capital reserves. The ratios of Expected Losses to capital reserves are very low, ranging from 0.003 for Emerald to

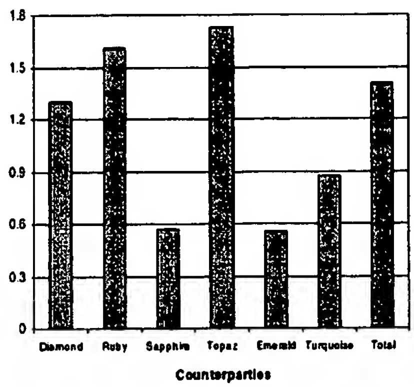

Figure 3: Ratio of MC/BIS exposure

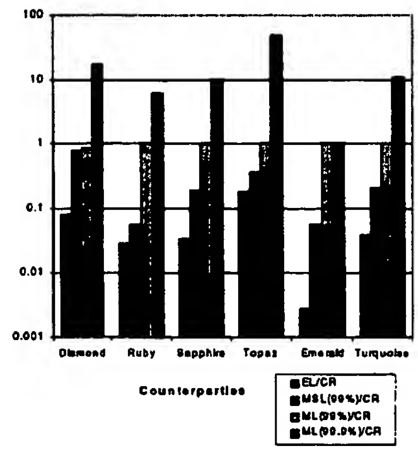

Figure 4: Ratios of MC/BIS losses

0.18 for Topaz. The ratios of *MSL(99%)* to the BIS capital reserves are all less than 1. The same is true for the ratios of *ML(99%)* to the BIS reserves. On the other hand, except for Emerald, the ratios of *ML(99.9%)* to the BIS reserves are very high, ranging from 6.25 for Ruby to 48.7 for Topaz.

The results indicate that credit reserves set to *ML(99%)* would be lower than the BIS capital reserves, while credit reserves set to *ML(99.9%)* would be much higher than the BIS capital reserves. Because the probability of a large credit loss for highly rated counterparties is very low, the first credit loss that is observed in a simulation typically falls between the 99% and 99.9% percentile. As a result, the decision taken on the appropriate confidence level significantly affects the amount of the reserve that should be set.

Next, we consider the capability of the methods to account for the impact of netting as a credit enhancement tool on credit exposures and losses. We compare the ratios of the Maximum Total Exposures with and without netting to the ratios of CEAs with and without netting.

Figure 5 shows that Spadina Bank can benefit significantly from netting agreements with counterparties with zero current AE, as evidenced by the ratios for Diamond, Sapphire, Emerald and Turquoise. In addition, for these four counterparties, the CEA ratios consistently exceed the *MTE(99%)* ratios, implying that the BIS approach does not adequately account for the effect of netting. On the other hand, for Ruby and Topaz, whose current Actual Exposures are positive, the BIS approach over-estimates the benefit of netting.

This counter-intuitive outcome is a result of the non-forward looking nature of the BIS approach: the Net-to-Gross Ratio is calculated based on current exposures only and does not consider future exposures. It is possible that portfolios have exposures that benefit less from netting in the future than today. While a Monte Carlo simulation over time captures the changing characteristics of the portfolio and the reduced netting benefits, the static BIS approach does not recognize any changes in the exposure profile over time and overestimates the risk reduction due to netting.

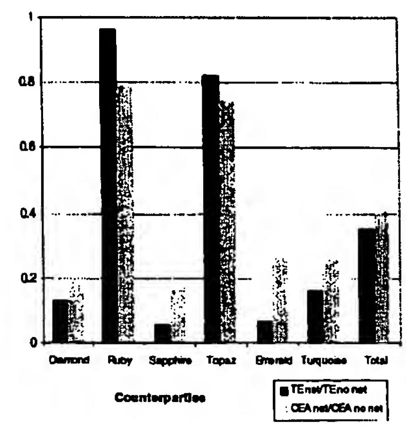

Figure 5: Netting/no netting exposure ratios

Conclusions

The amount of credit risk capital reserve has been a subject of much debate. Currently, the BIS method is the only approach banks can use to set their reserves. The results of this case study are sufficient to support many of the observations made by the ISDA evaluation of credit risk and regulatory capital requirements (ISDA 1998) concerning the deficiencies of the regulatory standard. On the other hand, a models-based approach, such as the Monte Carlo method used in this study, addresses many of these concerns.

The MC method can explicitly account for probabilities of default and recovery, and thus differentiate levels of credit risk among broad categories of credit. It estimates future exposure based on a set of market scenarios over an appropriate simulation horizon, thus accounting for the term structure of credit risk and the evolution of risk factors. This more accurate estimation of potential exposure leads to a more precise evaluation of the impact of netting and moneyness of the position. Moreover, the Monte Carlo simulation infrastructure provides a much richer set of information which credit managers can use to explain the causes of the exposure.

Because the BIS approach does not take into account the evolution of the exposure through time, its resulting reserve can be either too low or too high, depending on the nature of the transactions and the level of prudence required. Under the Monte Carlo method, the capital reserves are set to cover a maximum loss calculated at some level of confidence. The level chosen will significantly affect the amount of reserve that should be set.

We have explored the merits of the Monte Carlo based analysis of credit exposure and losses. The MC approach differentiates among categories of credit risk, accounts for the moneyness of the position and the changes in the values of the underlying risk factors, and recognizes the impact of the term structure of credit risk and the benefits of netting and of portfolio diversification. As well as specifically accounting for each of the factors that impacts credit exposure and loss, the simulation underlying the MC approach provides an explanation of the outcomes of the credit analysis.

However, some caution must be exercised when applying statistical analyses to any problem, and this case study provides excellent examples of these dangers. We discuss one illustrated by exposure measures and one by the losses.

Emerald's exposure profiles show the Expected Actual Exposure lying between the Maximum Actual Exposure profiles of *MAE(95%)* and *MAE(99%)*. The reason for this counter-intuitive result is that Actual Exposures are positive in less than 5% of the scenarios. Therefore, *MAE(95%)* is zero over time, while the Expected Actual Exposures are slightly positive. This illustrates the potential danger of using a percentile measure (such as Value-at-Risk) to measure credit risk.

The significant difference between the Maximum Losses, *ML(99.9%)* and *ML(99%)*, for all counterparties also highlights the danger of the use of a percentile measure. The lower the percentile, the greater the chance that the measure will miss outcomes that have low probabilities but extremely severe impacts.

Jeff Aziz joined Algorithmics in 1996 as a financial engineer. Currently, he is responsible for research and financial engineering related to consulting projects and provides financial engineering support to Client Services. Jeff has an MBA from McMaster University with a concentration in management science.

At the time this article was written, **Narat Charupat** was employed as a financial engineer focusing on credit risk functionality in RiskWatch and modeling new securities. Narat has a Ph.D. in Finance from York University and is now a professor in the Michael G. DeGroote School of Business at McMaster University.

References

Bank for International Settlements, 1997,
International Convergence of Capital Measurement and Capital Standards, Basle Committee on Banking Supervision, (July 1988, updated to April 1997)

Duffee, G., 1996,
"On measuring credit risks of derivative instruments," *Journal of Banking and Finance* 20(5): 805-833

Granadier, S. and B. Hall, 1995,
"Risk-based capital standards and the riskiness of bank portfolios: credit and factor risks," NBER Working Paper (#5178)

Hull, J. and A. White, 1995,
"The impact of default risk on the prices of options and other derivative securities," *Journal of Banking and Finance* 19(2): 299-322

Irving, R., 1997,
"The internal question," *Credit Risk: A Supplement to Risk Magazine* 10(7): 36-38

Jarrow, R. and S. Turnbull, 1995,
"Pricing derivatives on financial securities subject to credit risk," *Journal of Finance* 50(1): 53-85

Jamshidian, F. and Y. Zhu, 1997,
"Scenario simulation: theory and methodology," *Finance and Stochastics* 1(1): 43-67

Kreinin, A., 1997,
"Mean reverting processes in multi-step Monte Carlo simulation," Working Paper, Algorithmics Inc.

Kreinin, A., 1998,
"Distribution of economic losses under credit risk," Working Paper, Algorithmics, Inc.

Institute of International Finance, Inc.,
Report of the Working Group on Capital Adequacy: Recommendations for Revising the Regulatory Capital Rules for Credit Risk, March 1998

International Swaps and Derivatives Association,
Credit Risk and Regulatory Capital, March 1998

Standard & Poor's, Ratings and Performance 1997,
Stability and Transition, McGraw-Hill, New York: 1998

Endnotes

1. More generally, credit risk may also include the risk of loss that results from changes in expectations of the counterparty's likelihood to default, as reflected by an upgrade or downgrade of its credit rating.
2. Note that if the contract's value decreases monotonically with time, then Potential Exposure to that contract is zero.
3. Total Exposures are not discounted. As a result, the values are in the dollar value of the years that peak exposures occur.
4. The maximum values in the profiles of Expected Actual Exposures are not comparable to the tabulated values of the Maximum Total Exposure.

An Integrated Market and Credit Risk Portfolio Model

Ian Iscoe, Alex Kreinin and Dan Rosen

We present a multi-step model to measure portfolio credit risk that integrates exposure simulation and portfolio credit risk techniques. Thus, it overcomes the major limitation currently shared by portfolio models with derivatives. Specifically, the model is an improvement over current portfolio credit risk models in three main aspects. First, it defines explicitly the joint evolution of market factors and credit drivers over time. Second, it models directly stochastic exposures through simulation, as in counterparty credit exposure models. Finally, it extends the Merton model of default to multiple steps. The model is computationally efficient because it combines a Mark-to-Future framework of counterparty exposures and a conditional default probability framework.

Credit risk modeling is one of the most important topics in risk management and finance today. The last decade has seen the development of models for pricing credit risky instruments and derivatives, for assessing the credit worthiness of obligors, for managing exposures of derivatives and for computing portfolio credit losses for bonds and loan portfolios. In light of these financial innovations and modeling advances the Basle Committee on Banking Supervision has taken the first steps to amend current regulation and is reviewing the applicability of internal credit risk models for regulatory capital (Basle Committee on Banking Supervision, 1999a, 1999b).

However, common practice still treats market and credit risk separately. When measuring market risk, credit risk is commonly not taken into account; when measuring portfolio credit risk, the market is assumed to be constant. The two risks are then 'added' in *ad hoc* ways, resulting in an incomplete picture of risk.

There are two categories of credit risk measurement models: Counterparty Credit Exposure models and Portfolio Credit Risk models.

Derivative desks traditionally manage credit risk by monitoring and placing limits on counterparty credit exposures. **Counterparty exposure** is the economic loss that will be incurred on all outstanding transactions if a counterparty defaults, unadjusted by possible future recoveries. **Counterparty exposure models** measure and aggregate the exposures of all transactions with a given counterparty. In the BIS regulatory model, potential exposures are given by an add-on factor multiplying the notional of each transaction (Basle Committee on Banking Supervision, 1988). Although simple to implement, the model has been widely criticized because it does not accurately account for future exposures. Since exposures of derivatives such as swaps depend on the level of the market when default occurs, models must capture not only the actual exposure to a counterparty at the time of the analysis but also its potential future changes. Recently, more advanced methods based on Monte Carlo simulation (Aziz and Charupat 1998) have been implemented by financial institutions. By simulating counterparty portfolios through time over a wide range of scenarios, these models explicitly capture the contingency of the market on derivative portfolios and credit risk. Furthermore, they can accurately model natural offsets, netting, collateral and various mitigation techniques used in practice.

Since their main focus is on risk at the counterparty level, counterparty credit risk models do not generally attempt to capture portfolio effects such as the correlation between counterparty defaults. In contrast, **Portfolio Credit Risk (PCR) models** measure credit capital and are specifically designed to capture portfolio effects, specifically obligor correlations. They include CreditMetrics (JP Morgan 1997), CreditRisk+ (Credit Suisse Financial Products 1997), Credit Portfolio View (Wilson 1997a and 1997b) and KMV's Portfolio Manager (Kealhofer 1996). Although superficially they appear quite different—the models differ in their distributional assumptions, restrictions, calibration and solution—Gordy (1998) and Koyluoglu and Hickman (1998) show an underlying mathematical equivalence among these models. However, empirical work shows generally that all PCR models yield similar results if the input data is consistent (Crouhy and Mark 1998; Gordy 1998).

A major limitation of all current PCR models is the assumption that market risk factors, such as interest rates, are deterministic. Hence, they do not account for stochastic exposures. While this assumption has less consequence for portfolios of loans or floating rate instruments, it has great impact on derivatives such as swaps and options. Ultimately, a comprehensive framework requires the full integration of market and credit risk.

In this paper, we present a multi-step, stochastic model to measure portfolio credit risk that integrates exposure simulation and portfolio credit risk methods. Through the explicit modeling of stochastic exposures, the model overcomes the major limitation currently shared by portfolio models in accounting for the exposure caused by instruments with embedded derivatives. By combining a Mark-to-Future framework of counterparty exposures (see Aziz and Charupat 1998) and a conditional default probability framework (see Gordy 1998; Koyluoglu and Hickman 1998; Finger 1999), we minimize the number of scenarios where expensive portfolio valuations are calculated, and can apply advanced Monte Carlo or analytical techniques that take advantage of the problem structure.

We restrict this paper to a 'default mode' model; that is, the model measures credit losses arising exclusively from the event of default. However, default mode models cannot account for deals that have direct contingency on migrations (e.g., credit trigger features) without further modifications. Although perhaps computationally intensive, it is not difficult to extend the model to account for migration losses. Note, however, that since credit migrations are actually changes in expectations of future defaults, a multi-step model captures migration losses indirectly.

Specifically, the model presented in this paper is an improvement over current portfolio models in three main aspects:

- First, it defines explicitly the joint evolution of market risk factors and credit drivers. **Market factors** drive the prices of securities and **credit drivers** are non-idiosyncratic factors that drive the credit worthiness of

obligors in the portfolio. Factors are general and can be microeconomic, macroeconomic, economic and financial.

- Second, it models directly stochastic exposures through simulation, as do the Counterparty Credit Exposure models. In this sense, it constitutes an integration of counterparty exposure and Portfolio Credit Risk models.
- Finally, it extends the Merton model of default (1974), as used, for example, in CreditMetrics, to multiple steps. It explicitly solves for multi-step thresholds and conditional default probabilities in a general simulation setting.

The rest of the paper is organized as follows. We begin by introducing a general framework for Portfolio Credit Risk Models. The framework is first illustrated through the commonly known single-step model with deterministic exposures. Next, we present the multi-step, stochastic model in two stages. First, we extend the single-step model with deterministic exposures to account for stochastic exposures, and second, we extend that model to a multi-step version. The paper closes with some concluding remarks and an outline of future work.

Framework for Portfolio Credit Risk Models

Current portfolio models fit within a generalized underlying modeling framework. Gordy (1998) and Koyluoglu and Hickman (1998) first introduced the framework to facilitate the comparison between the various models. Finger (1999) further points out that formulating the models in this framework permits the use of powerful numerical tools known in probability that can improve computational performance by dramatically reducing the number of scenarios required. The main idea behind the framework is that conditional on a scenario all defaults and rating changes are independent. A state-of-the-world is a complete specification at a point in time of the relevant economic and financial credit drivers and market factors (macroeconomic, microeconomic, financial, industrial, etc.) that drive the model. A **scenario** is defined by a set of states-of-the-world over time. In a single-period model there is a direct correspondence between a state-of-the-world and a scenario; in a multi-period model a scenario corresponds to a path of states-of-the-world over time.

In this section, we introduce the basic components of the framework, which we subsequently use to present various models. We make several steps explicit in the framework, which were previously implicit in the original presentations. This further specification permits us to present the models in a manner that better explains the assumptions made and allows us to address the generalizations of the model.

The framework consists of five parts:

Part 1: Risk factors and scenarios. This is a model of the evolution of the relevant systemic risk factors over the analysis period. These factors may include both credit drivers and market factors.

Part 2: Joint default model. Default and migration probabilities vary as a result of changing economic conditions. An obligor's probabilities are conditioned on the scenario at each point in time. The relationship between its conditional probabilities and the scenario is obtained through an intermediate variable, called the obligor's credit worthiness index. Correlations among obligors are determined by the joint variation of conditional probabilities across scenarios.

Part 3: Obligor exposures, recoveries and losses in a scenario. The amount that will be lost if a credit event occurs (default or migration) as well as potential recoveries are computed under each scenario. Based on the level of the market factors in a scenario at each point in time, Mark-to-Future (MtF) exposures for each counterparty are obtained accounting for netting, mitigation and collateral. Similarly, recovery rates in the event of default can be state dependent.

Part 4: Conditional portfolio loss distribution in a scenario. Conditional upon a scenario, obligor defaults are independent. Various techniques based on the property of independence of obligor defaults can be applied to obtain the conditional portfolio loss distribution.

Part 5: Aggregation of losses in all scenarios. Finally, the unconditional distribution of portfolio credit losses is obtained by averaging the conditional loss distributions over all possible scenarios.

We illustrate the framework with a single time step Portfolio Credit Risk model with deterministic exposures, PCR_SD. Common notation and key concepts are also introduced.

Next, the model is extended to allow for stochastic exposures in a single-step setting, PCR_SS. Finally, we present a third model, PCR_MS, which allows for multiple time steps and stochastic exposures.

A set of four tables (Appendix 1) summarize the features of the models and highlight the similarities and differences of the models presented here. Table A1 presents a summary of the features of the three PCR models. Table A2 summarizes definitions of the risk factors and scenarios in Part 1 of the framework. Table A3 summarizes the components of the joint default model of Part 2. Table A4 summarizes the calculations for conditional obligor losses, conditional portfolio losses and unconditional losses of Parts 3 to 5 of the framework.

PCR_SD: Single-Step with Deterministic Exposures

The first model, PCR_SD, measures single-step portfolio credit losses with deterministic obligor exposures and recovery rates. This is a two-state form of the CreditMetrics model. We consider a default mode model, where default is driven by a Merton model.

Consider a portfolio with N obligors or accounts. Each obligor belongs to one of $N_s < N$ sectors. We assume that obligors in a sector are statistically identical. The grouping of obligors into sectors facilitates the estimation and solution of the problem.

Part 1. Risk Factors and Scenarios

Consider the single period $[t_0, t]$ where, generally, $t = 1$ year. In this single period model a scenario corresponds to a state-of-the-world. At the end of the horizon, t, the scenario is defined by q^c systemic factors, the credit drivers, which influence the credit worthiness of the obligors in the portfolio.

Denote by $x(t)$ the vector of factor returns at time t; i.e., $x(t)$ has components $x_k(t) = \ln(r_k(t)/r_k(t_0))$, where $r_k(t)$ is the value of the k-th factor at time t. Assume that at the horizon the returns are normally distributed: $x(t) \sim N(\mu, Q)$, where μ is a vector of mean returns and Q is a covariance matrix. Denote by $Z(t)$, the vector of normalized factor returns; i.e., $Z_k(t) = (x_k(t) - \mu_k) / \sigma_k$. For ease of exposition, and without loss of generality, assume that the factor returns are independent; independent factors can always be obtained, for example, by applying Principal Component Analysis to the original economic factors.

Part 2. Joint Default Model

The joint default model consists of three components. First, the definition of unconditional default probabilities. Second, the definition of a credit worthiness index for each obligor and the estimation of a multi-factor model that links the index to the credit drivers. Finally, a model of obligor default, which links the credit worthiness index to the probabilities of default, is used to obtain conditional default probabilities. Below, we explain these components in more detail.

Denote by τ_j the time of default of obligor j, and by $p_j(t)$ its **unconditional probability of default**, the probability of default of an obligor in sector j by time t:

(1) $$p_j(t) = Pr\{\tau_j \leq t\}$$

Note that all obligors in sector j have the same unconditional probability of default. We assume that unconditional probabilities for each sector are available from an internal model or from an external agency.

The **credit worthiness index**, Y_j, of obligor j determines the credit worthiness or financial health of that obligor at time t. Whether an obligor is in default can be determined by considering the value of its index. We assume that Y_j, a continuous variable, is related to the credit drivers through a linear, multi-factor model:

(2) $$Y_j(t) = \sum_{k=1}^{q^c} \beta_{jk} Z_k(t) + \sigma_j \varepsilon_j$$

where

$$\sigma_j = \sqrt{1 - \sum_{k=1}^{q^c} \beta_{jk}^2}$$

is the volatility of the idiosyncratic component associated with sector j, β_{jk} is the sensitivity of the index of obligor j to the k-th factor and ε_j, $j = 1, 2, \ldots, N$, are independent and identically distributed standard normal variables. Thus, the first term on the right side of Equation 2 is the systemic component of the index while the second term is the specific, or idiosyncratic, component. Note that the distribution of the index is standard normal; it has zero mean and unit variance.

Since all obligors in a sector are statistically identical, obligors in a given sector share the same multi-factor model. However, while all obligors in a sector share the same β_{jk} and σ_j, each has its own idiosyncratic, uncorrelated component, ε_j.

The **conditional probability of default** of an obligor in sector j, $p_j(t; Z)$, is the probability that an obligor in sector j defaults at time t, conditional on scenario Z:

(3) $$p_j(t;Z) = Pr\{\tau_j \leq t | Z(t)\}$$

The estimation of conditional probabilities requires a conditional default model which describes the functional relationship between the credit worthiness index Y_j (and hence the systemic factors) and the default probabilities p_j.

We assume that default is driven by a Merton model (Merton 1974). In the Merton model (Figure 1), default occurs when the assets of the firm fall below a given boundary or threshold, generally given by its liabilities. We consider that an obligor defaults when its credit worthiness index, Y_j, falls below a pre-specified threshold estimated from historical data. In this setting, an obligor's credit worthiness index, Y_j, can be interpreted as the standardized return of its asset levels. Default occurs when this index falls below α_j, the **unconditional default threshold**.

From an econometrics perspective, the Merton model is referred to as a probit model. It is conceptually straightforward to substitute a different default model, such as a logit model, as presented in Wilson (1997a, 1997b).

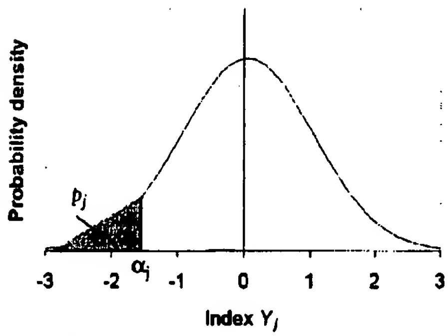

Figure 1: Merton model of default

The first step in the model defines the unconditional default threshold, α_j, for each obligor. The second step calculates the conditional default probabilities in each scenario.

The unconditional default probability of obligor j is given by

(4) $$p_j = Pr\{Y_j < \alpha_j\} = \Phi(\alpha_j)$$

where Φ denotes the normal cumulative density function. For simplicity, we have dropped the dependence on time, t, from the notation.Thus, the unconditional threshold, α_j, is obtained by the inverse of Equation 4:

(5) $$\alpha_j = \Phi^{-1}(p_j)$$

The conditional probability of default is then the probability that the credit worthiness index falls below the threshold in a given scenario:

(6) $$p_j(Z) = Pr\{Y_j < \alpha_j | Z\}$$

$$= Pr\left\{ \sum_{k=1}^{q^c} \beta_{jk} Z_k + \sigma_j \varepsilon_j < \alpha_j \mid Z \right\}$$

$$= Pr\left\{ \varepsilon_j < \frac{\alpha_j - \sum_{k=1}^{q^c} \beta_{jk} Z_k}{\sigma_j} \right\}$$

$$= \Phi\left(\frac{\alpha_j - \sum_{k=1}^{q^c} \beta_{jk} Z_k}{\sigma_j} \right)$$

$$= \Phi(\hat{\alpha}_j(Z))$$

The **conditional threshold**, $\alpha_j(Z)$, is the threshold that the idiosyncratic component of obligor j, ε_j, must fall below for default to occur in scenario Z.

Note that obligor credit worthiness index correlations are uniquely determined by the default model and the multi-factor model, which links the index to the credit driver returns. The correlations between obligor defaults are then obtained from the functional relationship between the index and the event of default, as determined by the Merton model. For example, the indices of obligors that belong to the same sector are perfectly correlated if their idiosyncratic component is zero.

Part 3. Obligor Exposures and Recoveries in a Scenario

Define the exposure to an obligor j at time t, V_j, as the amount that will be lost due to outstanding transactions with that obligor if default occurs, unadjusted for future recoveries. An important property of PCR_SD is the assumption that obligor exposure is deterministic, not scenario dependent: $V_j \neq f(Z)$.

The economic loss if obligor j defaults in any scenario is

(7) $$L_j(Z) = V_j \cdot (1 - \gamma_j)$$

where γ_j is the recovery rate, expressed as a fraction of the obligor exposure. Recovery, in the event of default, is also assumed to be deterministic. (Expressing the recovery amount as a fraction of the exposure value at default does not necessarily imply instantaneous recovery of a fraction of the exposure when default occurs.)

The distribution of conditional losses for each obligor is given by

$$L_j(Z) = \begin{cases} V_j \cdot (1-\gamma_j) & \text{with prob.} \quad p_j(Z) \\ 0 & \text{with prob.} \quad 1-p_j(Z) \end{cases} \quad (8)$$

Part 4. Conditional Portfolio Loss Distribution in a Scenario

Conditional on a scenario, Z, obligor defaults are independent. This follows from Equation 6 and the assumption that the idiosyncratic components of the indices are independent. To determine whether an obligor default occurs in a scenario, all that remains to be sampled is its idiosyncratic component.

In practice, the computation of conditional losses can be onerous. In the most general case, a Monte Carlo simulation can be applied to determine portfolio conditional losses. However, the observation that obligor defaults are independent permits the application of more effective computational tools. Some of these techniques are described in Credit Suisse (1997), Finger (1999) and Nagpal and Bahar (1999).

For the purpose of exposition only, consider a portfolio with a very large number of obligors, each with a small marginal contribution. In this case, we can use the Law of Large Numbers (LLN) to estimate conditional portfolio losses. As the number of obligors approaches infinity, the conditional loss distribution converges to the mean loss over that scenario; the conditional variance and higher moments become negligible. Hence, the conditional portfolio losses, $L(Z)$, are given by sum of the expected losses of each obligor:

$$L(Z) = \sum_{j=1}^{N} E(L_j(Z)) = \sum_{j=1}^{N} V_j \cdot (1-\gamma_j) \cdot p_j(Z) \quad (9)$$

Assuming that the LLN is appropriate simplifies the presentation which permits us to focus this discussion on the differences and similarities among the PCR_SD model and the stochastic and multi-step models that follow. Other methods include the application of the Central Limit Theorem (which assumes the number of obligors is large, but not necessarily as large as that required for the LLN), the application of moment generating functions with numerical integration or the application of probability generating functions with a discretization of exposures.

Part 5. Aggregation of Losses in All Scenarios

Unconditional portfolio losses are obtained by averaging the conditional losses over all scenarios. The distribution of unconditional portfolio losses is given by

$$Pr\{L_P < \lambda\} = \int_Z Pr\{L(Z) < \lambda\} dF(Z) \quad (10)$$

where L_P denotes the unconditional portfolio losses, λ denotes the level of losses and $F(Z)$ is the distribution of Z.

The aggregation of losses is generally obtained by performing a Monte Carlo simulation on the risk factor returns. Alternatively, analytical solutions are available under some restrictions (see Nagpal and Bahar (1999)).

The first column of Table A1 (Appendix 1) summarizes the features of the PCR_SD model. In the first columns of Tables A2 to A4, we summarize the components of the PCR_SD model associated with the framework.

PCR_SS: Single-Step Model with Stochastic Exposures

The second model developed in this paper relaxes the assumption of deterministic exposures and recoveries of the previous model. The PCR_SS model measures single-step portfolio credit losses due to default and assumes that obligor exposures and recoveries are stochastic.

Part 1. Risk Factors and Scenarios

The definition of risk factors and scenarios is similar to that in the PCR_SD model, with the difference that now a set of market factors are introduced which are determinant in the calculations of credit exposures.

Consider the single period $[t_0, t]$. At the end of the horizon, t, the scenario is now defined by q factors of which q^m are market factors and $q^c = q - q^m$ are credit drivers (this separation is for ease of exposition only and in no way restricts the model).

Denote by $x(t)$ the vector of factor returns at time t. In general, we use the superscript m to denote quantities related to market factors and the superscript c to denote those related to credit drivers. Thus, x^m and x^c denote the factor returns of the market factors and credit drivers, respectively. The first q^m components of $x(t)$ correspond to x^m and the following q^c components to x^c.

Assume that both credit driver and market factor returns are normally distributed: $x(t) \sim N(\mu, Q)$. Denote by $Z(t)$ the vector of normalized credit driver returns; i.e., $Z_k(t) = (x_k^c(t) - \mu_k^c(t))/\sigma_k(t)$. As in the PCR_SD model, assume that the components of Z are independent. Note that normalized, independent returns are required for

the credit drivers only; the returns of the market factors can follow more general models.

A scenario is described by an outcome of the returns vector x, or equivalently by a joint outcome of the vectors x^m and Z.

Part 2. Joint Default Model

The joint default model of PCR_SS is identical to that of PCR_SD. Conditional default probabilities are as given by Equation 6. In the PCR_SS model, the vector Z contains only the standardized returns of the credit drivers, $Z_k(t)$.

Part 3. Obligor Exposures and Recoveries in a Scenario

The main difference between the PCR_SS model and the previous deterministic model is that in this model obligor exposures are stochastic. The exposure to obligor j at time t, $V_j(x^m)$, varies by scenario as a function of the market risk factors; i.e., $V_j = f(x^m)$ (the index t is dropped for simplicity).

Exposure for each obligor is obtained through a single-step MtF simulation of all outstanding transactions, accounting for all netting agreements, mitigation and collateral. We refer to the table of obligor exposures by scenario, $V_j(x^m)$, as the Exposure MtF Table.

We further allow for recoveries in the event of default, γ_j, to be stochastic. The economic loss incurred if obligor j defaults in a given scenario is

$$L_j(x^m, Z) = V_j(x^m) \cdot (1 - \gamma_j(x^m, Z)) \qquad (11)$$

Because it is difficult to estimate the correlations, it is common to assume that recoveries are independent of the risk factors.

The table of obligor conditional losses by scenario, $L_j(x^m, Z)$ is referred to as the Obligor Losses MtF Table.

The loss distribution for each obligor is then

$$L_j(x^m, Z) = \begin{cases} V_j(x^m) \cdot \left(1 - \gamma_j(x^m, Z)\right) & \text{with prob. } p_j(Z) \\ 0 & \text{with prob. } 1 - p_j(Z) \end{cases} \qquad (12)$$

Part 4. Conditional Portfolio Loss Distribution in a Scenario

As in the PCR_SD model, obligor defaults are independent conditional on a scenario Z. The main difference, of course, is that the exposures and recoveries are now also a function of the scenario. Thus, if the portfolio contains a very large number of obligors, each with a small marginal contribution, the LLN dictates that conditional portfolio losses converge to the sum of the expected losses of each obligor:

$$L(x^m, Z) = \sum_{j=1}^{N} E\left\{L_j(x^m, Z)\right\} \qquad (13)$$

$$= \sum_{j=1}^{N} V_j(x^m) \cdot \left(1 - \gamma_j(x^m, Z)\right) \cdot p_j(Z)$$

Part 5. Aggregation of Losses in All Scenarios

Unconditional portfolio losses are obtained by averaging the conditional losses over all scenarios:

$$Pr\{L_P < \lambda\} = \int_{(x^m, Z)} Pr\left\{L(x^m, Z) < \lambda\right\} dF(x^m, Z) \qquad (14)$$

where $F(x^m, Z)$ is the joint distribution of the market risk factors and credit drivers. This integral is generally computed using a Monte Carlo simulation.

The second column of Table A1 (Appendix 1) summarizes the features of the PCR_SS model. In the second columns of Tables A2 to A4 we summarize the components of the PCR_SS model associated with the framework.

PCR_MS: Multiple-Step Model with Stochastic Exposures

The previous single-step, stochastic model, PCR_SS, is now extended to a multi-step setting. Model PCR_MS measures multi-step portfolio credit losses due to default and assumes that obligor exposures and recoveries are stochastic. Default is driven by a multi-step extension of a Merton model.

In this section, the full derivation of the discrete time model is presented. Appendix 2 introduces the problem of determining the credit worthiness process for the continuous time analog. The resolution of this problem will be addressed in future work.

Part 1. Risk Factors and Scenarios

Assume M multiple discrete time steps during the period $[t_0, T]$: $t_0 < t_1 < t_2 < \ldots < t_M = T$. A state-of-the-world at each time t_i is defined by a realization of q factors, out of which q^m are market factors and q^c credit drivers, respectively.

Denote by $r(t_i)$ the vector of risk factor values at time t_i and by $x(t_i)$ the vector of factor returns from time t_0 to t_i; i.e., $x(t_i)$ has components $x_k(t_i) = \ln(r_k(t_i)/r_k(t_0))$.

where $r_k(t_i)$ is the value of the k-th factor at time t_i. Again, x^m and x^c denote the factor returns of the market factors and credit drivers, respectively.

We assume that the evolution of the vector of risk factor values over time is determined by a set of stochastic differential equations:

$$dr_k(t) = \mu_k(r,t)dt + \sum_{l=1}^{q} \sigma_{kl}(r,t)dw_l, \quad k = 1, \ldots, q \tag{15}$$

where μ_k denotes the instantaneous drift of factor k and w_l, $l = 1, 2, \ldots, q$, are uncorrelated Wiener processes. The matrix $\sigma = (\sigma_{kl})$ is such that $\sigma\sigma^T$ forms an instantaneous covariance matrix. In general, the parameters of the stochastic differential equation can be functions of both time and risk factors.

More specifically, assume that the credit driver returns follow an arithmetic Brownian motion with constant coefficients:

$$dx_k^c(t) = \mu_k^c dt + \sum_{l=1}^{q^c} \sigma_{kl}^c dw_l^c, \quad k = 1, \ldots, q^c \tag{16}$$

No additional assumptions are made concerning the process for the market risk factors; the process for the market risk factors is as defined in Equation 15.

A scenario, in the discrete time setting, is then described by an outcome of the return vectors $x(t_i)$, $i = 1, \ldots, M$, (or, equivalently, by a joint outcome of the vectors r^m and r^c). Thus, in the multi-step model, a scenario is a path of states-of-the-world over time, i.e., a scenario is defined as $x = \{x(t_i), i = 1, \ldots, M\}$.

Part 2. Joint Default Model

The joint default model consists of three components: unconditional default probabilities and thresholds, a multi-factor model for the credit worthiness index of each obligor and a default model.

In the discrete time setting, the time of default of obligor j, τ_j, can take values t_i, $i = 1, \ldots, M$. The **unconditional probability of default** at time t_i, $p_j(t_i)$, is the probability that default of obligor j occurs in the i-th time step:

$$p_j(t_i) = Pr\{\tau_j = t_i\} \tag{17}$$

Denote by $P_j(t_n)$ the **unconditional cumulative probability of default** of an obligor in sector j by time t_n:

$$P_j(t_n) = Pr\{\tau_j \le t_n\} = \sum_{i=1}^{n} Pr\{\tau_j = t_i\} = \sum_{i=1}^{n} p_j(t_i) \tag{18}$$

Unconditional default probability term structures for each sector are an input to the model. They may be estimated from an internal model, from an external agency, or inferred from one period unconditional transition matrices, assuming a Markovian process.

In the development of the single-step models we noted that an obligor's credit worthiness index, Y_j, can be interpreted as the single-step standardized return of its asset levels. In the multi-step model the credit worthiness index of each obligor evolves through time. For a given obligor j, $A_j(t_i)$ is the level of the index. The return of the index up to time t_i is $y_j(t_i)$; i.e., $y_j(t_i) = \ln A_j(t_i)/A_j(t_0)$. Finally, $Y_j(t_i)$ is the standardized return on the index

$$Y_j(t_i) = \frac{y_j(t_i) - \mu_j(t_i)}{\sigma_j(t_i)} \tag{19}$$

where $\mu_j(t_i)$ and $\sigma_j(t_i)$ are respectively the mean and volatility of the index returns. $Y_j(t_i)$ has zero mean and unit volatility at every time step and is thus a canonical process. A canonical process is the process equivalent of a standard normal variable. In addition, we define the single period index returns:

$$\hat{y}_j(t_i) = \ln\left(\frac{A(t_i)}{A(t_{i-1})}\right) = y_j(t_i) - y_j(t_{i-1})$$

and

$$\hat{Y}_j(t_i) = \frac{\hat{y}_j(t_i) - \hat{\mu}_j(t_i)}{\hat{\sigma}_j(t_i)}$$

where $\hat{Y}_j(t_1) \equiv Y_j(t_1)$ and $\hat{\mu}_j(t_i)$ and $\hat{\sigma}_j(t_i)$ are the mean and volatility of the single period returns.

We assume that the index is related to the scenario through a continuous multi-factor model. The model for each obligor j is given by

$$dy_j(t) = \sum_{k=1}^{q^c} \beta_{jk}\, dx_k^c(t) + \sigma_j dw_j \tag{20}$$

where β_{jk} is the sensitivity of the index to the k-th credit driver; dw_j, $j = 1, \ldots, N$, are independent Wiener processes and σ_j is the volatility of the j-th idiosyncratic component. The canonical process $Y_j(t)$ can be derived from Equations 19 and 20. This process is the continuous analog to the single-step multi-factor model in Equation 2, except that the latter is standardized while the former is not.

In discrete time, the solution of Equation 20 can be written as

$$y_j(t_n) = \sum_{k=1}^{q^c} \beta_{jk}\, x_k^c(t_n) + \sigma_j \sum_{i=1}^{n} \sqrt{\Delta t_i}\,\varepsilon_i \tag{21}$$

where ε_i are independent and identically distributed standard normal variables and $\Delta t_i = t_i - t_{i-1}$. For ease of exposition we restrict attention to the case of uniform time steps, $\Delta t_i \equiv \Delta t$, $i = 1, 2, \ldots, M$.

Therefore, y_j is a stationary, independent-increments process and thus $\mu_j(t_i) = \mu_j(\Delta t)$,

$\sigma_j(t_i) = \sigma_j(\Delta t)$ and $\sigma_j(t_i) = \sigma_j(\Delta t)\sqrt{i}$ (since

$\sigma_j^2(t_i) = \sum_{k=1}^{i} \hat{\sigma}_j^2(t_k)$). Note that the random variables

$\{\tilde{\varepsilon}_{ji} = \hat{Y}_j(t_i), i = 0, 1, \ldots, M-1\}$ are independent, identically distributed and normal with zero mean and unit variance.

For the default model, we assume a multi-step Merton model. Default occurs in the first time step i that the index of the firm falls below the **unconditional default threshold**, $\alpha_{ji} = \alpha_j(t_i)$. The prescription of the model proceeds in two steps. The first step in the model defines the unconditional default thresholds for each obligor, α_{ji}. The second step calculates the conditional default probabilities in each scenario.

In the discrete, multi-step model, the time of default τ_j is the first time the credit worthiness index falls below the unconditional threshold:

$$\tau_j = \min_{i=1,\ldots M}\{t_i : Y_j(t_i) < \alpha_{ji}\}$$

Thus, the probability that default occurs in time step n is the probability that the index falls below the threshold in time step n and exceeds the threshold in each preceding period:

$$p_j(t_n) = Pr\{\tau_j = t_n\}$$

$$= Pr\{Y_j(t_1) > \alpha_{j1}, Y_j(t_2) > \alpha_{j2}, \ldots,$$

$$Y_j(t_{n-1}) > \alpha_{jn-1}, Y_j(t_n) < \alpha_{ji}\}$$

The calculation of the threshold for obligor j in the first time step, t_1, is similar to that of the single-step models. As given in Equation 4, the single-step unconditional probability of default is the cumulative normal of the unconditional threshold. More formally, we can also write the probability of default in the first time step as

$$p_j(t_1) = \Phi(\alpha_{j1}) = \int_{-\infty}^{\alpha_{j1}} \phi(v)dv \tag{22}$$

where

$$\phi(v) = \frac{e^{-v^2/2}}{\sqrt{2\pi}}$$

The threshold α_{j1} is the standard normal quantile associated with the unconditional probability appearing on the left side of Equation 22:

$$\alpha_{j1} = \Phi^{-1}\{p_j(t_1)\} \tag{23}$$

Note the similarity between Equation 23 and Equation 5.

The probability that default occurs in the second time step, t_2, can be written as

$$p_j(t_2) = Pr\{\tau = t_2\} = Pr\{Y_j(t_1) > \alpha_{j1}, Y_j(t_2) < \alpha_{j2}\} \tag{24}$$

$$= Pr\left\{Y_j(t_1) > \alpha_{j1}, \frac{\hat{Y}_j(t_1) + \hat{Y}_j(t_2)}{\sqrt{2}} < \alpha_{j2}\right\}$$

$$= Pr\left\{\tilde{\varepsilon}_{j1} > \alpha_{j1}, \tilde{\varepsilon}_{j1} + \tilde{\varepsilon}_{j2} < \tilde{\alpha}_{j2}\right\}$$

$$= \iint_{\substack{u > \alpha_1 \\ u+v < \tilde{\alpha}_{j2}}} \phi(u)\phi(v)du dv$$

$$= \int_{-\infty}^{\tilde{\alpha}_{j2}} \int_{\alpha_{j1}}^{\infty} \phi(u)\phi(v-u)du dv$$

where $\tilde{\alpha}_{j2} = \alpha_{j2}\sqrt{2}$. Default does not occur in the first time step, thus the limits of integration are $[\alpha_{j1}, \infty)$; default must occur in the second time step, thus the limits of integration are $(-\infty, \tilde{\alpha}_{j2}]$.

The unconditional threshold for the second time step, α_{j2}, is defined implicitly in Equation 24 from the probability $p_j(t_2)$ and the t_1-threshold α_{j1}.

More generally, for any time step n, the threshold α_{jn} is determined implicitly from the default probabilities and the thresholds at all previous time steps:

$$p_j(t_n) = Pr\{\tau = t_n\} \tag{25}$$

$$= \int_{-\infty}^{\tilde{\alpha}_{jn}} \int_{\tilde{\alpha}_{jn-1}}^{\infty} \cdots \int_{\tilde{\alpha}_{j1}}^{\infty} \phi(v_n - v_{n-1})\phi(v_{n-1} - v_{n-2})$$

$$\ldots \phi(v_2 - v_1)\phi(v_1)dv_1 \ldots dv_{n-1}dv_n$$

where $\tilde{\alpha}_{ji} = \alpha_{ji}\sqrt{i}$.

Given the cumulative default probability curve for each sector, thresholds can be computed using a Monte Carlo method which solves recursively for the limits of the integrals in Equation 25.

We define the **conditional probability of default**, $p_j(t_n;x^c)$, as the probability that default of an obligor in sector j occurs in the n-th time step conditional on the realization of the credit drivers up to time t_n:

$$p_j(t_n;x^c) = Pr\left\{\tau_j = t_n \middle| x^c(t_i), i = 1, \ldots, n\right\} \tag{26}$$

The computation of the conditional default probabilities is as follows: for the first time step, the conditional probability of default is, as in the previous models, given by

$$p_j(t_1;x^c) = Pr\left\{Y_j(t_1) < \alpha_{j1} \middle| x^c(t_1)\right\} \tag{27}$$

$$= Pr\left\{\frac{\left[\sum_{k=1}^{q^c} \beta_{jk}\, x_k^c(t_1) + \sigma_j\sqrt{\Delta t}\,\varepsilon_1\right] - \mu_j(t_1)}{\sigma_j(t_1)} < \alpha_{j1} \middle| x^c(t_1)\right\}$$

$$= Pr\left\{\varepsilon_1 < \frac{\sigma_j(t_1)\alpha_{j1} + \mu_j(t_1) - \sum_{k=1}^{q^c} \beta_{jk}\, x_{k1}^c}{\sigma_j\sqrt{\Delta t}}\right\}$$

$$= \Phi\left(\frac{\bar{\alpha}_{j1} - \sum_{k=1}^{q^c} \beta_{jk}\, x_{k1}^c}{\sigma_j\sqrt{\Delta t}}\right)$$

The threshold, adjusted by the drift and volatility of the index returns, is $\bar{\alpha}_{ji} = \sigma_j(t_i)\alpha_{ji} + \mu_j(t_i)$. Note that Equation 27 is equivalent to Equation 6.

For the second time step, the conditional probability is given by

$$p_j(t_2;x^c) = Pr\left\{Y_j(t_1) > \alpha_{j1}, Y_j(t_2) < \alpha_{j2} \middle| x^c(t_1), x^c(t_2)\right\}$$

$$= Pr\left\{\varepsilon_1 > \frac{\bar{\alpha}_{j1} - \sum_{k=1}^{q^c} \beta_{jk}\, x_{k1}^c}{\sigma_j\sqrt{\Delta t}}, (\varepsilon_1 + \varepsilon_2) < \frac{\bar{\alpha}_{j2} - \sum_{k=1}^{q^c} \beta_{jk}\, x_{k2}^c}{\sigma_j\sqrt{\Delta t}}\right\}$$

Then, in general, for time step n

$$p_j(t_n;x^c) = \left\{Pr \bigcap_{i=1}^{n-1} \{Y_j(t_i) > \alpha_{ji}\}, Y_j(t_n) < \alpha_{jn} \middle| x^c(t_i), i = 1, \ldots, n\right\} \tag{28}$$

where

$$u_{ji} = \frac{\bar{\alpha}_{ji} - \sum_{k=1}^{q^c} \beta_{jk}\, x_{ki}^c}{\sigma_j\sqrt{\Delta t}}$$

Equation 28 can be restated as

$$p_j(t_n;x^c) = Pr\left\{\bigcap_{i=1}^{n-1}\left[\sum_{l=1}^{i} \varepsilon_l > u_{ji}\right], \sum_{l=1}^{n} \varepsilon_l < u_{jn}\right\} \tag{29}$$

The right side of Equation 29 can be computed using numerical integration. Details of the computation of the multi-step conditional default probabilities are presented in Appendix 3.

Part 3. Obligor Exposures and Recoveries in a Scenario

As in the single-step stochastic model, PCR_SS, obligor exposures are stochastic. However, in this model, the exposure to obligor j is dependent on the path of the market risk factors up to time t_i, $V_{ji}(x^m) \equiv V_j(t_i) = f(x^m(t_k)), 0 \leq k \leq i$.. Since exposures at various times are summed, each is already discounted to today. Thus, discounted exposures express the capital that must be held today to cover future defaults (unadjusted for recoveries).

Exposures for each obligor are obtained through a multi-step simulation of all outstanding transactions, accounting for all netting agreements, mitigation and collateral. Aziz and Charupat (1998) present examples of the computation of these exposures. The table of obligor exposures over every scenario, $V_j(x^m, t_i)$ is referred to as the Multi-step Exposure MtF Table.

The economic loss if obligor j defaults in time step i, is the exposure of obligor j at time step i, net of recoveries, where it is also assumed that recoveries in the event of default, γ_{ji}, are stochastic:

$$L_{ji}(x) = V_{ji}(x^m) \cdot \left(1 - \gamma_{ji}(x^m, x^c)\right) \tag{30}$$

The probability of this default is $p_j(t_i,x^c)$. Thus, for every time step i, the distribution of conditional obligor losses is given by

$$L_{ji}(x) = \begin{cases} V_{ji}(x^m) \cdot \left(1 - \gamma_{ji}(x^m, x^c)\right) & \text{with prob. } p_j(t_i;x^c) \\ 0 & \text{with prob. } 1 - p_j(t_i;x^c) \end{cases} \tag{31}$$

The table of conditional obligor losses, $L_{ji}(x)$, is referred to as the Multi-step MtF Table of Obligor Losses.

Part 4. Conditional Portfolio Loss Distribution in a Scenario

At each time step, obligor defaults are independent, conditional on a scenario. The losses in a given scenario are simply the sum of the losses at each time step in that scenario.

If the portfolio contains a very large number of obligors, each with a small marginal contribution, the LLN dictates that conditional portfolio losses at each time step converge to the sum of the expected losses of each obligor:

$$\text{(32)} \qquad L(t_i, x) = \sum_{j=1}^{N} V_{ji}(x^m) \cdot \left(1 - \gamma_{ji}(x^m, x^c)\right) \cdot p_j\left(t_i; x^c\right)$$

Expected portfolio losses in a given scenario are the sum of the expected losses in each time step:

$$\text{(33)} \qquad L(x) = \sum_{i=1}^{M} L(t_i, x)$$

Part 5. Aggregation of Losses in All Scenarios

Unconditional portfolio losses are obtained by averaging the conditional losses over all scenarios:

$$\text{(34)} \qquad Pr\{L_p < \lambda\} = \int_x Pr\{L(x) < \lambda\} dF(x)$$

where $F(x)$ is the probability distribution in the scenario space. This integral is generally computed using Monte Carlo techniques.

The third column of Table A1 (Appendix 1) summarizes the features of the PCR_MS model. In the third columns of Tables A2 to A4 we summarize the components of the PCR_MS model associated with the framework.

Concluding Remarks

We have presented a new multi-step Portfolio Credit Risk model that integrates exposure simulation and advanced portfolio credit risk methods. The integrated model, PCR_MS, overcomes a major limitation currently shared by portfolio models in accounting for the credit risk of portfolios whose exposures depend on the level of the market.

Specifically, the model presented in this paper is an improvement over current portfolio models in three main aspects. First, it defines explicitly the joint evolution over time of market risk factors and credit drivers. Second, it models directly stochastic exposures through simulation, as in Counterparty Credit Exposure models. Finally, it extends the Merton model of default to multiple steps. Although the latter seems conceptually straightforward, the resulting mathematical model is not trivial. Moreover, expressing the model so that it is amenable to an efficient solution is essential.

The model is computationally efficient because it combines a Mark-to-Future (MtF) framework of counterparty exposures and a conditional default probability framework. The computational benefits are threefold:

- First, the number of scenarios for which expensive portfolio valuations are made is minimized.
- Second, the model is based on the same computations used for monitoring counterparty exposures and placing limits at the desks. A MtF framework allows users to exploit these computations and use them for both counterparty exposures and portfolio credit risk. This is not only important for computational purposes, but also leads to more consistent enterprise risk measurement.
- Third, advanced Monte Carlo or analytical techniques that take advantage of the problem structure can be used to solve the problem faster and more accurately than standard Monte Carlo methods.

We have restricted this paper to a default mode model. It is not conceptually difficult to extend the model to account for migration losses as well. Note that since credit migrations are simply changes in expectations of future defaults, a multi-step model partially and indirectly captures migration losses. Future work will address these issues in detail. ®

Ian Iscoe joined Algorithmics in 1996. He is the technical manager, Risk Analytics and he is responsible for software development related to scenario generation, dynamic trading strategies, optimization and RiskMetrics. Ian has a Ph.D. in Mathematics from Carleton University.

Alexander Kreinin joined Algorithmics in 1995. Currently, he works as a senior mathematician in the Research Group and is the manager of Quantitative Research. Alex has a Ph.D. in Probability and Statistics from Vilnius University in Lithuania.

Dan Rosen joined Algorithmics in 1995. He is currently the director of Research and is responsible for Algorithmics' financial and mathematical research and joint projects with academic institutions. Dan has a Ph.D. in Applied Sciences from the Department of Chemical Engineering and Applied Chemistry at the University of Toronto.

References

Aziz, J. and N. Charupat, 1998, "Calculating credit exposure and credit loss: a case study," *Algo Research Quarterly*, 1(1): 31-46.

Basle Committee on Banking Supervision, 1988, *International Convergence of Capital Measurements and Capital Standards*, July, http://www.bis.org

Basle Committee on Banking Supervision, 1999a, *Credit Risk Modeling: Current Practices and Applications*, Basle, April, http://www.bis.org

Basle Committee on Banking Supervision, 1999b, *A New Capital Adequacy Framework*, Consultative Paper, Basle, June, http://www.bis.org

CreditMetrics: The Benchmark for Understanding Credit Risk, Technical Document, 1997, New York, NY: JP Morgan Inc.

Credite Suisse Financial Products, 1997, *CreditRisk+: A Credit Risk Management Framework*, New York, NY.

Crouhy, M. and R. Mark, "A comparative analysis of current credit risk models," paper presented at the conference Credit Modeling and Regulatory Implications, London, September 1998.

Finger, C., 1999, "Conditional approaches for CreditMetrics portfolio distributions," *CreditMetrics Monitor*, April: 14-33.

Gordy, M., 1998, "A comparative anatomy of credit risk models," Federal Reserve Board, Finance and Economics Discussion Series, 1998-47.

Kealhofer, S., 1996, "Managing default risk in portfolios of derivatives," in *Derivative Credit Risk*, London: Risk Publications.

Koyluoglu, H.U. and A. Hickman, 1998, "Reconcilable differences," *Risk*, 11(10): 56-62.

Merton, R.C., 1974, "On the pricing of corporate debt: the risk structure of interest rates," *Journal of Finance*, 29: 449-470.

Nagpal K.M. and R. Bahar, 1999, "An analytical approach for credit risk analysis under correlated defaults," *CreditMetrics Monitor*, April: 51-74.

Wilson, T., 1997a, "Portfolio credit risk I," *Risk*, 10(9): 111-117.

Wilson, T., 1997b, "Portfolio credit risk II," *Risk*, 10(10): 56-61.

Appendix 1. Model Summary Tables
This appendix contains tables summarizing the three models presented.

Time steps	single-step	single-step	multi-step
Exposures and recoveries	deterministic	stochastic	stochastic
Credit events	default	default	default
Default model	Merton	Merton	Merton
Simulation factors	credit drivers	credit drivers market factors	credit drivers market factors
Factor distributions	credit drivers: standardized normal returns	credit drivers: standardized normal returns market factors: normal returns	credit drivers: normal returns market factors: general
Input data	(unconditional) one year default probabilities for each obligor/sector unconditional threshold	(unconditional) one year default probabilities for each obligor/sector unconditional threshold	term structure of (unconditional) default probabilities for each obligor/sector unconditional thresholds
Output results	Exposure MtF Table Obligor Losses MtF Table	Exposure MtF Table Obligor Losses MtF Table	Multi-step Exposure MtF Table Multi-step MtF Table of Obligor Losses

Table A1: Feature summary of Portfolio Credit Risk models

Credit risk factors	credit driver returns; standardized normal $Z_k(t) \sim N(0,1)$	credit driver returns; standardized normal $Z_k(t) \sim N(0,1)$	Equation 16 credit driver returns; normal $dx_k^c(t) = \mu_k^c dt + \sum_{l=1}^{q} \sigma_{kl}^c dw_l^c$
Market risk factors	not applicable	market factor returns; normal $x^m(t) \sim N(\mu, Q)$	Equation 15 market factor returns; general $dr_k^m(t) = \mu_k(r,t)dt + \sum_{l=1}^{q} \sigma_{kl}(r,t)dw_l$
Scenarios	single-step	single-step	multi-step

Table A2: Part 1—Definition of risk factors and scenarios

Unconditional default probability	**Equation 1** $p_j(t) = \Pr\{\tau_j \le t\}$	**Equation 1** $p_j(t) = \Pr\{\tau_j \le t\}$	**Equation 17** $p_j(t_i) = \Pr\{\tau_j = t_i\}$ **Equation18** $P_j(t_n) = \Pr\{\tau_j \le t_n\}$
Credit worthiness index	**Equation 2** $Y_j(t) = \sum_{k=1}^{q^c} \beta_{jk} Z_k(t) + \sigma_j \varepsilon_j$ **where** $\sigma_j = \sqrt{1 - \sum_{k=1}^{q^c} \beta_{jk}^2}$	**Equation 2** $Y_j(t) = \sum_{k=1}^{q^c} \beta_{jk} Z_k(t) + \sigma_j \varepsilon_j$ **where** $\sigma_j = \sqrt{1 - \sum_{k=1}^{q^c} \beta_{jk}^2}$	**Equation 21** $y_j(t_n) = \sum_{k=1}^{q^c} \beta_{jk} x_k^c(t_n) + \sigma_j \sum_{i=1}^{n} \sqrt{\Delta t}\, \varepsilon_i$
Unconditional default threshold	**Equation 5** $\alpha_j = \Phi^{-1}(p_j)$	**Equation 5** $\alpha_j = \Phi^{-1}(p_j)$	**Equation 25, solve for** $\alpha_{jn} = \bar{\alpha}_{jn} / \sqrt{n}$ $p_j(t_n) = \int_{-\infty}^{\bar{\alpha}_{jn}} \int_{\bar{\alpha}_{jn-1}}^{\infty} \cdots \int_{\bar{\alpha}_{j1}}^{\infty} \phi(v_n - v_{n-1}) \phi(v_{n-1} - v_{n-2}) \ldots \phi(v_2 - v_1) \phi(v_1)\, dv_1 \ldots dv_{n-1} dv_n$
Conditional default probabilities	**Equation 3** $p_j(t;Z) = \Pr\{\tau_j \le t \mid Z(t)\}$ **Equation 6** $p_j(Z) = \Phi\left(\frac{\alpha_j - \sum_{k=1}^{q^c} \beta_{jk} Z_k}{\sigma_j}\right)$	**Equation 3** $p_j(t;Z) = \Pr\{\tau_j \le t \mid Z(t)\}$ **Equation 6** $p_j(Z) = \Phi\left(\frac{\alpha_j - \sum_{k=1}^{q^c} \beta_{jk} Z_k}{\sigma_j}\right)$	**Equation 26** $p_j(t_n; x^c) = \Pr\left\{\tau_j = t_n \mid x^c(t_i), i = 1, \ldots, n\right\}$ **Equation 29** $p_j\left(t_n; x^c\right) = \Pr\left\{\bigcap_{i=1}^{n-1}\left[\sum_{l=1}^{i} \varepsilon_l > u_{ji}\right], \sum_{l=1}^{n} \varepsilon_l < u_{jn}\right\}$ **where** $u_{ji} = \bar{\alpha}_{ji} - \left(\sum_{k=1}^{q^c} \beta_{jk} x_{ki}^c\right) / \sigma_j \sqrt{\Delta t}$

Table A2: Part 2—Joint default models

Obligor exposure	V_j	$V_j(x^m)$	$V_{ji}(x^m)$
Obligor recovery	γ_j	$\gamma_j(x^m, Z)$	$\gamma_{ji}(x^m, x^c)$
Obligor losses in the event of default	Equation 7 $L_j(Z) = V_j \cdot (1-\gamma_j)$	Equation 11 $L_j(x^m, Z) = V_j(x^m) \cdot (1-\gamma_j(x^m, Z))$	Equation 30 $L_{ji}(x) = V_{ji}(x^m) \cdot \left(1-\gamma_{ji}(x^m, x^c)\right)$
Conditional portfolio losses (LLN)	Equation 9 $L(Z) = \sum_{j=1}^{N} E\{L_j(Z)\} = \sum_{j=1}^{N} V_j \cdot (1-\gamma_j) \cdot p_j(Z)$	Equation 13 $L(x^m, Z) = \sum_{j=1}^{N} V_j(x^m) \cdot \left(1-\gamma_j(x^m, Z)\right) \cdot p_j(Z)$	Equation 32 $L(t_i, x) = \sum_{j=1}^{N} V_{ji}(x^m) \cdot \left(1-\gamma_{ji}(x^m, x^c)\right) \cdot p_j\left(t_i; x^c\right)$ Equation 33 $L(x) = \sum_{i=1}^{M} L(t_i, x)$
Unconditional portfolio credit losses	Equation 10 $Pr\{L_p < \lambda\} = \int_Z Pr\{L(Z) < \lambda\} dF(Z)$	Equation 14 $Pr\{L_p < \lambda\} = \int_{(x^m, Z)} Pr\left\{L(x^m, Z) < \lambda\right\} dF(x^m, Z)$	Equation 34 $Pr\{L_p < \lambda\} = \int_x Pr\{L(x) < \lambda\} dF(x)$

Table A4: Parts 3, 4 and 5—Obligor exposure, recovery and losses

Appendix 2. The Canonical Credit Worthiness Process

Consider a single obligor. The canonical credit worthiness process, $Y(t)$, is described by

$$Y(t) = \frac{w_t}{\sqrt{t}}$$

where w_t is a one-dimensional Wiener process. In the continuous time Merton model, the time of default τ is the first time when the process falls below a boundary $\alpha(t)$:

$$\tau = inf_t\{Y(t) < \alpha(t)\}$$

Let the function $P(t)$, with $0 \le P(t) \le 1$ for all $t<T$ and $\frac{dP}{dt} \ge 0$ represent the probability of default before time t; i.e.,

$$P(t) = Pr\{\tau < t\}$$

This continuous time problem is depicted in Figure A1.

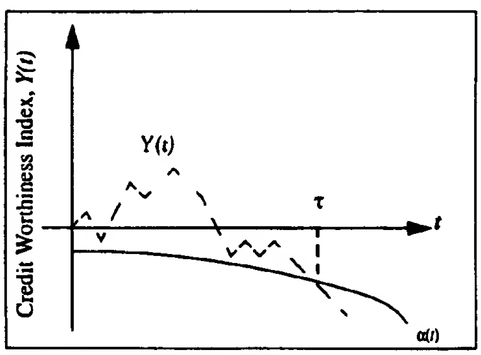

Figure A1: Continuous time default model

Appendix 3. Computation of Multi-Step Conditional Default Probabilities

In general, for time step n, the conditional default probabilities of a given obligor are given by Equations 28 and 29:

$$p_j\left(t_n; x^c\right) = Pr\left\{\bigcap_{i=1}^{n-1}\{Y_{ji} > \alpha_{ji}\}, Y_{jn} < \alpha_{jn} \middle| x^c(t_i), i = 1, \ldots, n\right\}$$

$$= Pr\left\{\bigcap_{i=1}^{n-1}\left[\sum_{l=1}^{i} \varepsilon_l > u_{ji}\right], \sum_{l=1}^{n} \varepsilon_l < u_{jn}\right\}$$

where

$$u_{ji} = \left(\bar{\alpha}_{ji} - \sum_{k=1}^{q} \beta_{jk} x_{ki}\right) / (\sigma_j \sqrt{\Delta t})$$

For simplicity, the index j, denoting a given obligor, is removed from the notation.

Denote by

$$A_n = \left\{\bigcap_{i=1}^{n-1}\left\{\sum_{l=1}^{i} \varepsilon_l > u_i\right\}, \sum_{l=1}^{n} \varepsilon_l < u_n\right\}$$

and

$$B_n = \bigcap_{i=1}^{n}\left\{\sum_{l=1}^{i} \varepsilon_l > u_i\right\}$$

Then it follows that

$$Pr(B_{n-1}) = Pr(A_n) + Pr(B_n)$$

The probability $Pr(B_n)$ is a function of n variables, $G_n(u_1, \ldots, u_n)$. The function $G_n(u_1, \ldots, u_n)$ satisfies the relation

(A1) $$G_{n+1}(u_1, \ldots, u_{n+1}) = G_n(u_1, \ldots, u_n)\Psi(u_{n+1} - u_n) + \int_{u_n}^{\infty} G_n(u_1, u_2, \ldots, u_{n-1}, u)\phi(u_{n+1} - u)du$$

where $\Psi(t) = \bar{\Phi}(t) = Pr\{\varepsilon_{n+1} > t\}$.

The integrals in Equation A1 are evaluated using numerical integration techniques.

Liability Using Dynamic Portfolio

Many institutions, especially large corporations and governments, incur liabilities by issuing debt in order to raise capital. This capital is usually raised in both short and long-term bond markets. In raising this capital, it is necessary to consider the risk/cost trade-off between issuing short term debt, which implies frequent rebalancing at uncertain interest rates, and long term bonds, which typically pay a higher coupon. In liability management, **cost** may be defined as the expected value of cumulative interest expenses and **risk** as the variability of cumulative interest expenses.

The liability management problem is characterized by long time horizons, typically several years. Over these lengthy horizons "portfolio effects" such as issuance, aging, coupon payments and maturities must be considered. To model the impact of these effects the positions must change over time. **Dynamic portfolios** are portfolios with positions that change over time according to **dynamic portfolio strategies** — decision rules that define how the positions change. These rules may be a simple deterministic, pre-specified rebalancing schedule or a more complicated set of rules that depend on future market events and/or portfolio characteristics.

Standard implementations of common risk management tools, such as RiskMetrics Value at Risk (VaR) (JP Morgan 1996), assume there are no portfolio effects. In fact, RiskMetrics VaR analysis is based on portfolios in which the positions do not change over the horizon of the analysis. While this assumption may be reasonable for the time horizon of traditional risk management (10 days, say), it is untenable for longer horizons. If these portfolio effects are unaccounted for, the estimates of future portfolio characteristics may be biased under longer time horizons.

Management Strategies

Richard Black and Colin Telmer

Moreover, VaR methodologies base risk measures on 'marked-to-market' values. Although suitable for many applications, the VaR methodology is not as appropriate when stakeholders of an organization are concerned about factors other than market value. For example, shareholders, rightly or wrongly, may consider price, earnings and profit indicators. In the context of an institution issuing debt, the market value is perceived to be less important, and managers typically look at the development in debt costs, as defined by some function of interest payments.

This paper couples dynamic portfolio strategies with scenario generation in a simulation methodology to determine how the interest expenses of a liability portfolio evolve over time. The scenario generation methodology is assumed to forecast risk-factor innovations that may have non-zero means. Another standard assumption in risk management is that these factors have mean zero; this assumption leads to biased results in an analysis that extends over a long time horizon. The reader is referred to Kim, Malz and Mina (1999) for a discussion of the issues associated with forecasting financial variables over a long time horizon.

The introduction of dynamic strategies enables all the tools of risk management, including Monte Carlo simulation, VaR and stress testing, to be applied to asset/liability management problems. Precise valuation models that account for portfolio effects improve the accuracy of the calculations. Conversely, it also allows the tools of asset/liability management, for example, risk/return analysis, to be applied to risk management problems.

The usefulness of dynamic portfolio strategies is illustrated through an extended example. In the example, we consider the problem of designing a government's debt program. This problem can be viewed as a constrained optimization problem in which an asset/liability manager must minimize risk subject to a cost constraint and also, in this case, a constraint that the government maintains a given cash balance in its account.

The example is organized as follows. First, the problem is described in detail to illustrate how a simple dynamic portfolio strategy can be used to meet the government's objective under a particular nominal interest rate scenario. In this case, 3 month bills and 5 year bonds are used by the government to maintain a given cash balance. The results of this analysis are stress tested. Then, in order to develop a better understanding of the risk/cost trade-off, the analysis is repeated for different dynamic portfolio strategies and cost levels. The result is an efficient frontier. Since there are only two types of instruments, this efficient frontier is simply a locus of risk/cost points and all strategies are efficient. To generate a more interesting set of results, the problem is enriched by the introduction of 2-year bonds. The set of dynamic portfolios on the new efficient frontier strictly dominates any others. Finally, a more complicated dynamic portfolio strategy is considered and its effect is discussed in terms of the risk/cost trade-off.

A priori, the example is expected to illustrate the conventional wisdom that short-term debt is cheap but risky and long-term debt is expensive but safe. Indeed, the results strongly support this. However, even this rather simple analysis suggests some shortcomings with this intuition. In particular, the results suggest that issuing too much long-term debt can actually increase risk. The increased risk associated with financing the higher interest payments of the long term debt more than offsets any reduction in risk incurred by issuing long-term debt.

Although asset/liability issues can be addressed using other techniques, dynamic portfolio strategies offer a

more intuitive and easily implemented approach that is capable of dealing with practical problems in a real-world setting. Before launching into the example, the next section provides a brief overview of other methods of addressing asset/liability management problems and how they relate to dynamic portfolio strategies.

Dynamic Portfolios

Although not yet prevalent in risk management practices, the history of dynamic portfolios is extensive. In its simplest form, single-period portfolio selection can be thought of as a dynamic portfolio problem. Traditional Markowitz mean-variance analysis (1952,1987) can be used to determine the optimal holdings in a portfolio of risky assets such that the optimal, rebalanced portfolio yields the minimum variance for a given rate of return. However, this approach allows for only a single rebalancing within the time horizon examined.

More interesting multi-stage dynamic portfolio problems may be modeled as the outcome of a dynamic stochastic programming problem. **Dynamic stochastic programming**, the study of procedures for decision making over time in a stochastic environment, deals well with uncertainty; however, the problem size escalates dramatically even for very small numbers of financial assets, time periods, possible return outcome values and constraints the government wishes to impose. Eppen and Fama (1968, 1971) model two- and three-asset problems using this technique and Daellenbach and Archer (1969) extend their work to include one liability. These models consider uncertainty of return and are dynamic, but only problems with a very small number of financial instruments can be analyzed simultaneously; hence, they are of limited use in practice.

Wolf (1969), Bradley and Crane (1972, 1980) and Lane and Hutchinson (1980) use stochastic decision tree models. Bradley and Crane apply their dynamic stochastic model to bond portfolio management. Their model, while useful for small problems, again becomes computationally unwieldy with even a few periods and possible outcomes. Kusy and Ziemba (1986) discuss a stochastic linear program under the uncertainty approach and compare it with Bradley and Crane's models. They show by simulation that the stochastic linear programming approach is superior to the decision tree dynamic programming approach developed by Bradley and Crane; however, the Kusy and Ziemba model does not account for final period effects nor is it truly dynamic since it is solved two periods at a time in a rolling fashion.

The Russell-Yasuda Kasai model, developed for the Yasuda Fire and Marine Insurance Co., Ltd. and described in Carino and Ziemba (1998) and Carino, Myers and Ziemba (1998), builds on this previous research to develop a large scale dynamic model with possibly dependent scenarios, final period effects, and all the relevant institutional and policy constraints of Yasuda Kasai's business enterprise. Although the Russell-Yasuda Kasai model is one of the first genuine commercial applications of dynamic stochastic programming, the complexity of the model limits its extensibility and adaptability by other institutions.

An alternative to modeling dynamic portfolios using stochastic dynamic programming is to use decision rules. These rules have several advantages over other methods: they are simple to communicate, they can truly capture the nature of the firm's behaviour and they are robust in the face of uncertainty. However, they may not be optimal in the sense that the dynamic portfolio resulting from a set of decision rules may or may not replicate the portfolio resulting from a stochastic dynamic programming exercise. Decision rules can vary dramatically from straightforward rules such as a 'buy-and-hold' rule to very complex rules involving derivative securities.

In the context of debt issuance, debt portfolio decision rules may vary widely in complexity. For instance, small institutions with access to well-developed capital markets may have the flexibility of choosing to issue debt at any maturity as their participation may have little impact on the overall market. In this case, their decision rules may be solely based on the trade-off between risk and cost. One possible decision rule, though not necessarily the optimal one, may simply be to issue debt at the lowest cost maturity along the term structure.

In contrast, the issuance behaviour of larger institutions may affect the overall market. Thus, their decision rules may have to account for additional criteria and therefore may be more complicated. For example, due to liquidity constraints, a large institution may not be able to issue all its debt at a single maturity, as a small institution can, but may have to spread the issuance across the term structure. Even so, it may not always be able to issue as much debt as required without affecting the debt price. Taken together, these problems imply that the decision rules for a large institution must account for both the concentration of issuance at a particular maturity as well as the absolute amount of issuance.

The decision rules used in this paper range from simple rules that do not depend on any future events to more complex, dynamic rules that depend on future events as they unfold. The complex rules are easily derived by modifying the simple decision rules. These rules are applied to the debt issuance problem introduced in the next section.

Debt Issuance

An example of government debt issuance is used to illustrate the benefits of dynamic portfolios. A stylized program of debt issuance is as follows. Each year the government's budget determines expected borrowing requirements net of any maturities of outstanding debt. Actual requirements are met by issuing bonds and bills using a bond issuance program that is pre-announced at the beginning of the horizon and a bill issuance program that dynamically adjusts to meet outstanding borrowing requirements. Any changes in borrowing requirements arising from, for example, changes in fiscal policy or changes in interest rates, are accommodated by an offsetting adjustment to the bill program.

The size of the pre-announced bond program relative to overall borrowing requirements is determined by considering the risks and costs inherent in the overall borrowing strategy. In this example, cost is measured by interest expenses cumulated over the entire horizon and risk is measured by the variability of the cumulative interest expenses. These definitions suggest that in an environment where the yield curve slopes, on average, upward, issuing short-term debt will often be cheaper than issuing long-term debt. However, short-term debt must be refinanced each period and as the interest rate at which this will occur is not known in advance, the variability in interest payments—which determines risk—associated with issuing short-term debt is higher than that associated with issuing long-term debt.

The value of the cash account today is equal to the value of the cash carried over from the previous period, plus any issuance in bonds and bills, less the value of interest payments and any settlements into the cash account from bonds and bills that mature in this period:

$$cash_t = cash_{t-1} + bills_t + bonds_t - maturities_t - interest_t \quad (1)$$

where $cash_t$ represents the government's cash account at time t, $cash_{t-1}$ represents the cash account from the previous period, $bills_t$ and $bonds_t$ represent new bill and bond issuance, respectively, and $maturities_t$ and $interest_t$ represent principal repayments and interest expenses.

If future capital requirements are known in advance and interest rates are deterministic, then all the terms of Equation 1 are known and the equation can be solved for total issuance at each time t $(bills_t+bonds_t)$. Designing a debt issuance program would be straightforward—a pre-announced schedule for both bonds and bills could be designed to meet the necessary requirements.

In a more realistic environment where capital requirements are not necessarily known in advance and interest rates are stochastic, the problem becomes more complicated and the program must be designed to accommodate future changes in these variables. In this case, given forecasts of future borrowing requirements and interest rates, Equation 1 can still be used to find an expected issuance program, but the program will depend on the actual (unknown) path of future interest rates and capital requirements.

To illustrate how dynamic portfolio strategies may be used in this debt issuance environment, an extended example is presented that begins with a very simple set of dynamic portfolio strategies.

A Simple Example

Following the stylized debt issuance program outlined above, the first example in this paper uses a simple set of dynamic portfolio strategies to develop a new debt issuance program. This issuance program is developed using a given yield curve referred to as the nominal scenario; this nominal scenario also serves as a reference curve for the sets of scenarios generated for stress analyses in later sections.

At the beginning of the horizon, the government's outstanding debt portfolio consists solely of treasury bills. The government initiates a bond program in an effort to reduce its interest cost and risk exposure. In light of this, the government wishes to analyze the risk and cost implications of this new program over a five-year horizon. The debt program in this example may be summarized as follows:

- Each quarter the government has a capital requirement that is met by issuing short- and long-term debt using treasury bills and bonds. A quarterly capital requirement exists to maintain a cash balance of $100,000.
- Part of this requirement is met by issuing bonds according to a pre-announced quarterly schedule. Each quarter, $5,000 worth of 5-year, fixed-rate coupon bonds are issued at par.
- The remainder of the requirement is raised using treasury bills. These are issued in a dynamic fashion such that in each quarter enough bills are issued to meet the overall capital requirement. In particular, sufficient bills are issued to cover any interest expenses and maturities in excess of new bond issuance. Bills are assumed to be 3-month bills (zero-coupon bonds).
- At the beginning of the five-year horizon, the institution has previously issued enough 3-month bills so that its current cash balance is $100,000.

From Equation 1, the assumption that the cash balance remains constant implies that interest payments evolve according to

$$interest_t = bills_t + bonds_t - maturities_t$$

or equivalently, if maturities are identified as being either bond or bill maturities

(2) $interest_t = net\ bill\ issuance_t + net\ bond\ issuance_t$

Figures 1 and 2 illustrate the evolution of the bill and bond program in an environment with an upward sloping yield curve. At the beginning of the horizon, there is $100,000 of outstanding 3-month bills that mature at the end of the first quarter. At that time, the principal repayment and interest due on these bills must be refinanced to maintain the required cash balance. Bonds equalling $5,000 are issued and the remaining refinancing requirement is met using bills. This can be seen in Figure 1. In the first quarter, bonds with a net worth of $5,000 are issued, interest expenses are approximately $1,000, and, using Equation 2, net bill issuance is $1,000 - $5,000 = -$4,000. In other words, $100,000 worth of matured bills and approximately $1,000 in interest are refinanced using $5,000 worth of bonds and approximately $96,000 worth of bills. After the first quarter, as illustrated in Figure 1, the bond program continues to issue $5,000 per quarter and net bill issuance remains negative. Over time, as shown in Figure 2, the higher bond issuance compared with bill issuance leads bonds to dominate the composition of government debt. Note also from Figure 2 that total debt increases over time as interest payments are financed through new issuance.

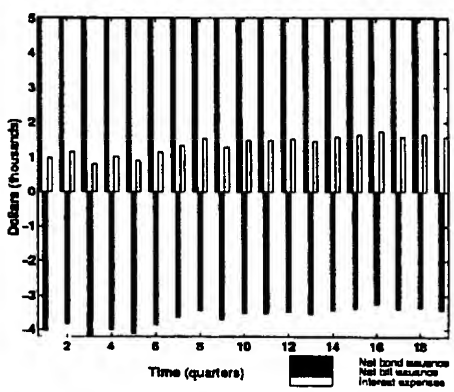

Figure 1: Future issuance patterns

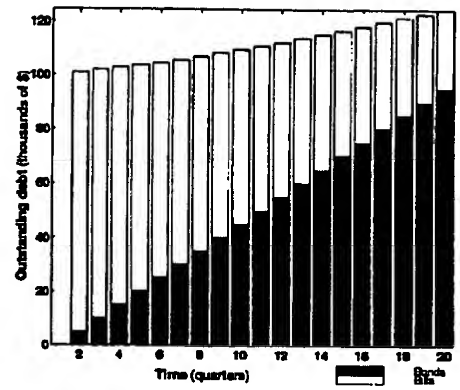

Figure 2: Future outstanding debt

Figures 3(a) and 3(b) graph interest expenses on a quarterly and cumulative basis. Again, it can be seen that interest expenses are increasing. More interestingly, it can also be seen that interest expenses vary over time. The interest curve evolves through the forward rates and because the interest curve is upward sloping, future spot rates change. This affects both the cost of raising funds using bills and, through its affect on the coupon rate, the cost of raising funds using bonds. As the *gross* issuance of bills is greater than the *gross* issuance of bonds, interest expenses due to bill issuance tend to be more volatile over time. (As these are 3-month bills, gross bill issuance is equal to the amount of bills outstanding which exceeds the $5,000 of gross bond issuance—see Figure 2).

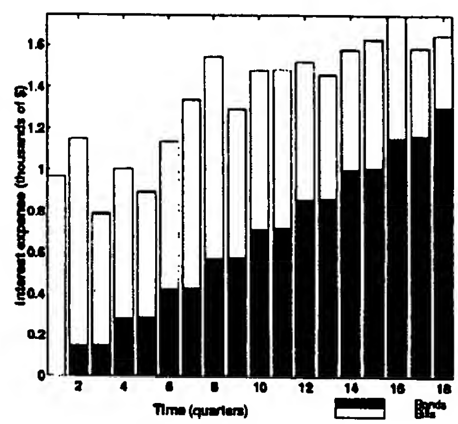

(a) Quarterly

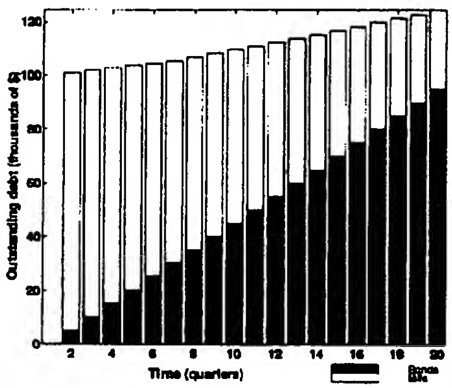

(b) Cumulative

Figure 3: Interest expense

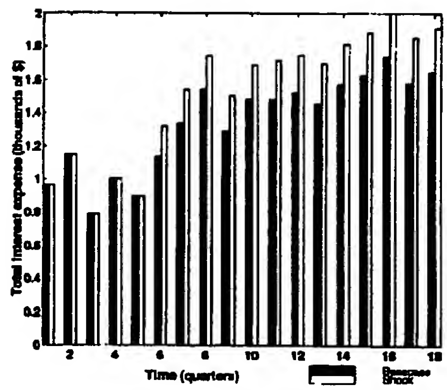

(a) Total interest expense

Alternative Issuance Patterns

To explore the impact on cost of various issuance patterns, the previous analysis is repeated with alternative issuance patterns. The alternative patterns range from a strategy of issuing no bonds (implying that all refinancing is done with 3-month bills) to one that calls for issuing $10,000 bonds, in increments of $2,500. Table 1 summarizes the results.

0	25,182
2,500	26,339
5,000	27,496
7,500	28,653
10,000	29,810

Table 1: Alternative issuance patterns

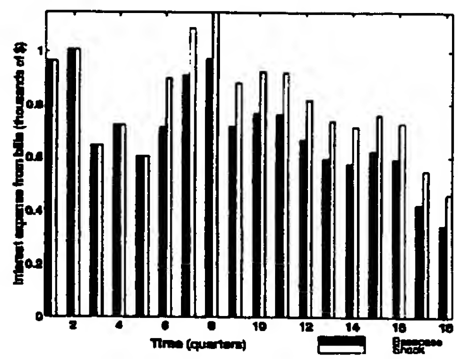

(b) Interest expense from bills

Recall that an upward sloping yield curve is used to calculate the interest exposure. It is apparent from Table 1 that portfolios with a higher bond issuance have higher cost, measured in cumulative interest payments. However, what is not apparent from examining the table is how the risk, measured by the volatility of these interest payments, changes with the amount of bond issuance. The risk associated with different issuance patterns is explored, first with a simple stress test and then in a Monte Carlo simulation.

A Simple Stress Test

The first simple test stresses the portfolio to assess the sensitivity of cumulative interest expense to the nominal interest rate scenario. The stress test is a 1% parallel shift in the yield curve that occurs in the second year of the planning horizon. Figures 4(a), 4(b), and 4(c) present the results corresponding to a $5,000 bond program.

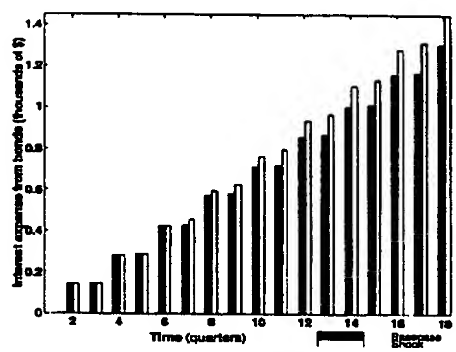

(c) Interest expense from bonds

Figure 4: Parallel shift in interest rates

Not surprisingly, the increase in interest rates leads to an increase in interest expenses. What is more relevant is that most of the increase in interest expenses, at least in the near term, is due to interest payments associated with the bill program. Recall that each quarter the entire bill program is 'rolled over' (that is, each quarter all the bills mature and are then reissued). Thus, the higher interest rates are felt immediately. The interest expenses associated with the bond program are less sensitive to the curve shift because there is less new bond issuance. Note, however, that interest expenses associated with the bond program increase over time as more bonds are issued at the higher rates.

Table 2 presents the results for the range of issuance strategies presented in Table 1. Portfolios with a higher bond issuance are less affected by the increase in interest rates. Again, portfolios comprising more bills will incur higher costs as the bills are rolled over in the higher interest rate environment.

[redacted]

0	25,182	29,846	4,665
2,500	26,339	30,450	4,111
5,000	27,496	31,053	3,557
7,500	28,653	31,656	3,001
10,000	29,810	32,269	2,449

Table 2: Interest rate stress test

The results of the stress test suggest that issuing more bonds leads to less variability in cumulative interest payments, albeit at a higher cost.

An Efficient Frontier

In this section we examine the risk associated with different issuance patterns in a Monte Carlo setting. The issuance patterns are those used in the previous stress test. The innovations to the nominal interest rate scenario are generated using a standard two-factor affine yield model similar to a model first proposed in Vasicek (1977). The model was calibrated using empirical moments estimated from the McCulloch and Kwon (1993) dataset.

Table 3 presents the results. Again, the mean cumulative interest expenses (the cost of the issuance patterns) increase with the amount of bond issuance. In contrast, however, the standard deviation (variability) of interest charges (the risk of the issuance pattern) decreases with the amount of bond issuance. This happens because long-term rates are less volatile than short-term rates and, more importantly, because there is less rollover risk associated with a large bond program.

[redacted]

0	22,680	4,238
2,500	24,551	3,523
5,000	26,521	2,896
7,500	28,492	2,425
10,000	30,463	2,215

Table 3: Risk and cost of alternative issuance patterns

Figure 5 displays the results for a large number of bond programs. Point A corresponds to a portfolio with a \$10,000 bond program and point B to a portfolio with a \$0 bond program. Points along the line connecting point A and point B represent portfolios with both bonds and bills in some combination. Moving from point A to point B corresponds to reducing the size of the bond program. As the bond issuance decreases, risk increases and cost decreases.

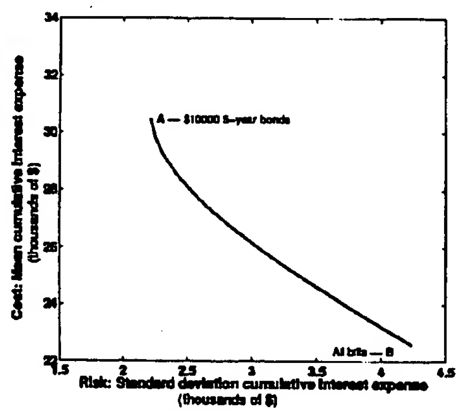

Figure 5: A simple efficient frontier

Figure 5 is a simple example of an **efficient frontier**—the locus of risk and cost combinations offered by portfolios of risky assets that yield the minimum risk for a given cost.

The efficient frontier may be used to determine the appropriate issuance rule for an institution given its risk appetite. Given the amount of risk that the organization wishes to bear, the bond program that the institution should employ may be found by vertically mapping the chosen risk level to the efficient frontier. This will lead to a bond program with the lowest possible expected cost for this level of risk; the actual cost may be found by horizontally mapping the point on the frontier to the vertical axis.

The efficient frontier in Figure 5 is simply a line—there are no interior points. This is because there are only two instruments available; in the next section, following the introduction of a third instrument, the frontier becomes a region in the risk/cost space.

An Increased Opportunity Set

In this section, 2-year, fixed-rate coupon bonds may be issued in addition to 3-month bills and 5-year bonds. As with 5-year bonds, the amount of issuance may vary between $0 and $10,000.

Figure 6 presents the results of the exercise superimposed on the previous results. The line that joins point A (all 5-year bonds) and point B (all 3-month bills) corresponds to the efficient frontier of the previous section. Point C corresponds to a portfolio of only 2-year bonds and point D to a portfolio of $10,000 worth of 2-year bonds and $10,000 worth of 5-year bonds. The feasible region is defined by the lines that connect point B with point C, point C with point D and point D with point B. Note that point A lies inside the feasible region.

Adding 2-year bonds to a portfolio of bills only (point B), moves the portfolio toward point C. The line BC is below BA, implying that, at the margin, the cost of reducing risk using 2-year bonds is less than the cost of using 5-year bonds. When issuing 2-year bonds rather than 5-year bonds, the increase in risk is more than offset by the decrease in cost.

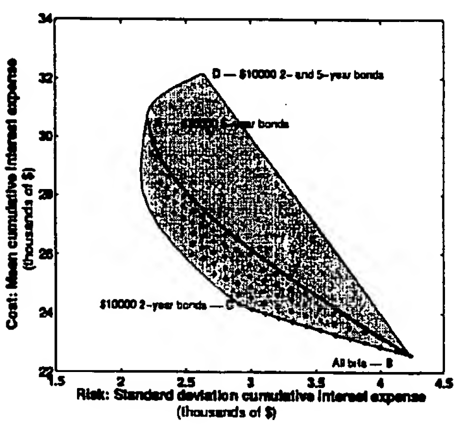

Figure 6: A more complicated efficient frontier

Moving along the frontier, from point C toward point D, 5-year bonds are added to a portfolio that comprises $10,000 worth of 2-year bonds. At the margin, this addition reduces risk and increases cost—up to a point. Eventually, the bond program becomes so large that risk starts to increase because the larger bond program generates more cash than required and thus, bill issuance is actually negative. Risk increases as the magnitude of the bill issuance becomes increasingly more negative.

As can be seen, the new efficient frontier dominates the old frontier everywhere—strategies that include 2-year bonds have both reduced risk and cost over any strategy that uses only bills and 5-year bonds. The portfolios that lie between the original frontier and the new frontier comprise bills and bonds, both 2-year and 5-year. The portfolios in the region bounded by ABC (labeled with a '*') correspond to portfolios that have a total bond program of less than $10,000. The portfolios in the region between line AC and the new efficient frontier represent bond programs of more than $10,000. The portfolios worth less than $10,000 have less risk because of the different risk/cost profile of 2-year bonds compared to 5-year bonds; portfolios worth more than $10,000 have less risk because programs with larger bond issuance tend to reduce risk simply because bonds are less risky.

Note that these results depend on the specification of the interest-rate scenarios. In general, yields and the variance/covariance of yields along the curve will determine the relative risks and costs of different strategies.

Note also that the simple, predetermined issuance strategies considered so far can result in negative bill issuance. The bill program is conditional on interest rates—if interest rates are lower than expected, then, depending on the predetermined bond program, the amount of bill issuance will be lower and may become negative. Constraints can be added to the problem that restrict the size of bond issuance such that bill issuance is always positive. Alternatively, conditional decision rules can be used to address this problem.

Conditional Decision Rules

In the previous section, the bond program is defined by a simple rule—a fixed amount of bonds is issued in each quarter for the duration of the planning horizon. Such simple rules are not very intuitive—institutions regularly update their bond programs in light of their current portfolio and market conditions. Such rules are also limiting in the sense that, by using more complicated issuance rules, portfolios that have less cost and less risk may be constructed.

Many institutions adjust the amount of bond issuance each quarter to maintain a pre-specified ratio of bonds to total outstanding debt. More specifically, if the actual proportion of bonds to total debt outstanding is greater than a target ratio, bond issuance is decreased and bill issuance is increased by an offsetting amount; if the proportion is less than the target the reverse happens—bond issuance is increased and bill issuance is decreased by an offsetting amount. The more complex issuance rules

considered in this section are based on this practice. The target ratio of outstanding bonds to total outstanding debt varies between zero and one for the different issuance functions examined. As well, the issuance of any one bond must be positive and may not exceed $10,000; thus, the results are comparable to the previous programs. Given these rules, in particular the limit on the maximum issuance size, the target ratio may not always be obtained immediately, especially if the target ratio is quite different from the current ratio.

Figure 7 presents the results for this set of strategies superimposed on the results in Figure 6. Point E corresponds to a portfolio of $10,000 5-year bonds and $10,000 2-year bonds. Any relaxation of the predetermined issuance rules will lead to an efficient frontier that lies on or outside that depicted in Figure 6. As can be seen, the new rule allows for a reduction in both risk and cost relative to the simple rules of predetermined issuance. More specifically, the more complicated rules allow for a reduction in risk and cost for portfolios with large bond issuance.

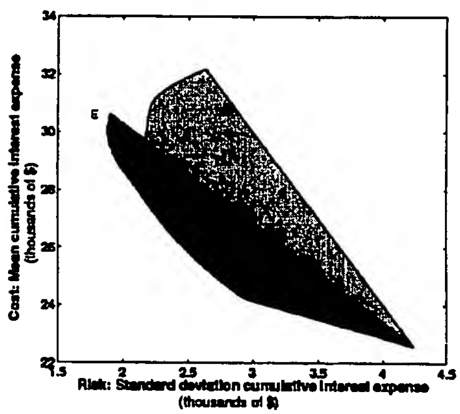

Figure 7: Efficient frontier with a variable bond program

Conclusions

Dynamic portfolio strategies give risk managers the ability to realistically forecast the cost/risk relationship inherent in portfolios over longer horizons. By examining longer horizons than those analyzed in traditional risk management, managers can better assess long-term risk including risk measured on a cash flow basis as well as on a market or fair value basis. Thus, a long-term Value-at-Risk can be calculated for a dynamic portfolio in a Mark-to-Future framework.

In the paper, the problem of debt issuance in a stochastic interest rate environment is used to illustrate how dynamic portfolio strategies may be applied. The paper specifies a number of example dynamic portfolio strategies, beginning with a very simple deterministic strategy based on a pre-specified rebalancing schedule applied to a portfolio of two instruments—a 3-month bill and a 5-year bond. The addition of a 2-year bond to the portfolio demonstrates the benefits of expanding the universe of instruments. Finally, the complexity of the dynamic strategy is increased by specifying a rule to dynamically adjust the amount of bond issuance each quarter in order to target a pre-specified ratio of fixed-rate debt to total outstanding debt. This additional flexibility allows further reductions in both cost and risk.

The task of constructing dynamic portfolio strategies for actual applications will likely result in strategies that are more complicated than those presented in this paper. Indeed, the design of new strategies is an important area of future research. A strategy of particular interest to risk managers might encapsulate limit-based risk management by comparing a VaR measure to a limit, permitting trading only when the VaR does not exceed the limit. Thus, dynamic strategies applied to risk management policies may translate into a risk/cost trade-off for the firm.◎

At the time this article was written, **Richard Black** headed up the Risk Management Unit at the Bank of Canada. In this role, he monitored the Bank's risk exposures and worked on debt strategy issues. Prior to that, he spent five years working on macroeconomic modeling within the Bank's Research Department. Richard has a B.A.(Hons) in Mathematics from the University of Canterbury in New Zealand and an M.A. in Economics from Simon Fraser University.

Colin Telmer joined the Government of Canada's Department of Finance in February 1998. As a financial economist in the Debt Management section, Colin works on debt strategy simulation model development. Colin has an M.A. in Economics and is currently working on a Ph.D. in Economics from Queen's University.

References

Bradley, S.P. and D.B. Crane, 1972,
"A dynamic model for bond portfolio management,"
Management Science 19: 139-151.

Bradley, S.P. and D.B. Crane, 1980,
"Managing a Bank Bond Portfolio Over Time,"
in *Stochastic Programming*, M.A.H. Dempster (ed.), London: Academic Press, 449-471.

Carino, D.R. and W.T. Ziemba, 1998,
"Formulation of the Russell-Yasuda Kasai financial planning model," *Operations Research*, 46(4): 433-449.

Carino, D.R., D.H. Myers and W.T. Ziemba, 1998,
"Concepts, technical issues, and uses of the Russell-Yasuda Kasai financial planning model,"
Operations Research, 46(4): 450-462.

Daellenbach, H.G. and S.A. Archer, 1969,
"The optimal bank liquidity: a multi-period stochastic model,"
Journal of Quantitative Analytics.

Eppen, G.D. and E.F. Fama, 1968,
"Solutions for cash balance and simple dynamic portfolio problems," *Journal of Business*, 41: 94-112.

Eppen, G.D. and E.F. Fama, 1971,
"Three asset cash balance and dynamic portfolio problems,"
Management Science, 17: 311-319.

RiskMetrics Technical Document, 4th Ed., 1996,
New York: JP Morgan Inc.

Kim, J., A.M. Malz and J. Mina, 1999,
"LongRun," New York: RiskMetrics Group.

Kusy, M. I. and W.T. Ziemba, 1986,
"A bank asset and liability management model,"
Operations Research, 34: 356-376.

Lane, M. and P. Hutchinson, 1980,
"A model for managing a certificate of deposit portfolio under uncertainty," in *Stochastic Programming: Proceedings of the First International Conference on Stochastic Programming*, Oxford: Academic Press.

Markowitz, H.M., 1952,
"Portfolio selection," *Journal of Finance*, 7(1): 77–91.

Markowitz, H.M., 1987,
Mean-Variance Analysis in Portfolio Choice and Capital Markets, Oxford: Basil Blackwell Ltd.

McCulloch, J.H. and H.C. Kwon, 1993,
"U. S. term structure data, 1947–1991,"
Ohio State University, Working Paper, 93-6.

Vasicek, O., 1977,
"An equilibrium characterization of the term structure,"
Journal of Financial Economics, 5(2): 177-188.

Wolf, C.R., 1969,
"A model for selecting commercial bank government security portfolios," *Review of Economics and Statistics*, 1(1): 40-52.

Beyond VaR

Financial institutions worldwide have devoted much effort to developing enterprise-wide systems that integrate financial information across their organizations to *measure* their institution's risk. Probabilistic measures, such as Value-at-Risk (VaR), are now widely accepted by both financial institutions and regulators for assigning risk capital and monitoring risk. Since development efforts have been driven largely by regulatory and internal requirements to report risk numbers, tools needed to *understand* and *manage* risk across the enterprise have generally lagged behind those designed to measure it.

Measuring risk is a *passive* activity; simply knowing one's VaR does not provide much guidance for managing risk. In contrast, risk management is a *dynamic* endeavor and it requires tools that help identify and reduce the sources of risk. These tools should lead to an effective utilization of the wealth of financial products available in the markets to obtain the desired risk profiles.

To achieve this, a comprehensive risk manager's toolkit must provide the ability to

- represent complex portfolios simply
- decompose risk by asset and/or risk factor
- understand how new trades affect the portfolio risk
- understand the impact of instruments' non-linearities and of non-normal risk factor distributions on portfolio risks
- understand complex, non-intuitive, market views implicit in the portfolio as well as in the investment policy or market liquidity
- generate potential hedges and optimize portfolios.

Robert Litterman (1996a, 1996b, 1997a, 1997b) recently described a comprehensive set of analytical risk management tools extending some of the insights originally developed by Markowitz (1952) and Sharpe (1964). Developed in close collaboration with the late Fisher Black and his colleagues at Goldman Sachs, these tools are based on a linear approximation of the portfolio to measure its risk and assume a joint (log)normal distribution of the underlying market risk factors, similar to the RiskMetrics VaR methodology (J.P. Morgan 1996). Litterman further emphasized the dangers of managing risk using only such linear approximations. However, in spite of their onerous assumptions, the insights provided by these tools are very powerful and hence constitute a solid basis for a risk management toolkit. (The reader is also referred to the related papers by Mark Garman (1996, 1997) on marginal VaR and risk decomposition.)

This is the first of a series of papers in which we present an extended *simulation-based risk management toolkit* developed on top of the analytical tools presented by Litterman. Simulation-based tools provide additional insights when the portfolio contains non-linearities, when the market distributions are not normal or when there are multiple horizons. In particular, these tools should prove very useful not only for market risk, but also for credit risk, where the exposure and loss distributions are generally skewed and far from normal. We further demonstrate that simulation-based tools can be used, sometimes even more efficiently, with other risk measures in addition to VaR. Indeed, they also uncover limitations of VaR as a coherent risk measure, as has been demonstrated by Artzner et al. (1998).

Simulation-based methods to measure VaR (historical or Monte Carlo) are generally much more computationally intensive than parametric methods (such as the delta-normal method popularized by RiskMetrics). Advances in computational simulation methods and hardware have rendered these methods practical for enterprise-wide risk measurement. However, it is widely believed that risk management tools based on simulation are impractical since they require substantial additional computational work (Dowd 1998). We demonstrate that efficient computational methods are available which generally require little or no additional simulation to obtain risk management analytics.

In this paper, we extend *marginal VaR analysis* to a simulation-based environment, compare the method with the parametric approach and apply it to two practical examples. We demonstrate how one can efficiently obtain the changes in the portfolio VaR, as measured by a simulation, that result from adding a small amount (or percentage) of an asset to the portfolio. We show that, since VaR is a homogeneous function of the positions, we can obtain an additive portfolio decomposition based on marginal VaR, as in the parametric case. We also investigate the trade risk profiles of a single asset or a class of assets. We discuss the properties of these tools, the errors that arise due to sampling, their limitations and possible extensions.

This paper is organized as follows. We first review parametric VaR and use its associated risk management tools to analyze a portfolio of foreign exchange forwards. We then derive the simulation-based tools and discuss the potential effects of sampling error. To demonstrate the methodology, we re-examine the foreign exchange portfolio (obtaining results consistent with the parametric version) and also consider a portfolio of stock options for which the parametric approach is inappropriate. We conclude by suggesting directions for further study.

Parametric VaR

The parametric, or delta-normal, method for calculating VaR assumes the existence of a set of market risk factors whose log price changes are joint normally distributed with zero mean; that is, if r_k is the log return on risk factor k, then $r \sim N(0, Q^*)$, where Q^* is the covariance matrix of risk factor returns. Consider a portfolio

composed of positions x, where x_i is the size of the holding in instrument i, for $i = 1, 2, \ldots, N$. As shown in the Appendix, the portfolio's $100(1-\alpha)\%$ VaR (which we denote $VaR(x)$, implicitly recognizing its dependence on α) is

$$VaR(x) = \sqrt{m(x)^T Q m(x)} \quad (1)$$

where $Q = Z_\alpha^2 Q^*$ is a scaled covariance matrix (Z_α denotes the standard normal z-value that delimits a probability of α in the right tail) and $m(x)$, known as the **VaR map** of the portfolio, is a vector of the portfolio's exposure to the risk factors. The elements of $m(x)$ equal the monetary value of the portfolio's position in each risk factor. Thus, the VaR map provides a reduced, or simplified, view of the portfolio from a risk management perspective. Note that

$$m(x) = \sum_{i=1}^{N} m^i x_i \quad (2)$$

where vector m^i is the VaR map of one unit of the i-th instrument (i.e., m_k^i is the exposure to risk factor k that results from holding a single unit of instrument i). Equation 2 shows that the portfolio VaR map is the sum of the instruments' VaR maps, weighted by position.

Trade Risk Profile and Best Hedge Position

Knowledge of how VaR changes with position size is critical for effective risk management. If we plot the portfolio VaR against the size of the position in a given instrument i (all other positions being fixed), we obtain the **trade risk profile** (TRP). As shown in Figure 1, the TRP has a unique minimum, which occurs at the best hedge position, x_i^{bh}. The best hedge position can be found analytically as described in the Appendix.

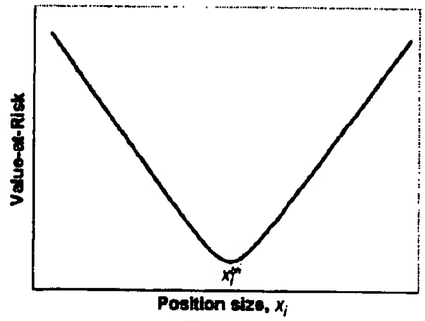

Figure 1: Trade risk profile

Marginal VaR

Managing risk requires an understanding of how new trades affect the portfolio risk. Thus, let us now consider the calculation of the marginal VaR, which measures the impacts of small changes in risk factor exposures or instrument positions on the portfolio VaR. From Equation 1, we find that the VaR gradient with respect to the risk factor exposures is

$$\nabla_m VaR(x) = \frac{Qm(x)}{VaR(x)} \quad (3)$$

The k-th element of $\nabla_m VaR(x)$ is the change in VaR that results from increasing the portfolio's exposure to the k-th risk factor (i.e., $m_k(x)$) by a single monetary unit.

Since the VaR map of the portfolio is the sum of the VaR maps for the positions (Equation 2), it follows that the derivative of VaR with respect to the i-th position is

$$\frac{\partial VaR(x)}{\partial x_i} = (m^i)^T (\nabla_m VaR(x)) \quad (4)$$

Equation 4 indicates the change in VaR due to adding one unit of instrument i to the portfolio (if $x_i < 0$, then this corresponds to reducing the short position). Note that this is simply the derivative of the trade risk profile for instrument i at the current position x_i.

CAD/USD .73 100d	CAD	100	0.73	0.5	2.5
CAD/USD .74 30d	CAD	30	0.74	1.0	-8.3
DEM/USD .57 60d	DEM	60	0.57	6.0	73.2
DEM/USD .59 120d	DEM	120	0.59	5.0	-28.2
FRF/USD .16 40d	FRF	40	0.16	8.0	83.3
JPY/USD .0091 11d	JPY	11	0.0091	10.0	-0.9

Table 1: FX portfolio

VaR Contribution

By decomposing VaR, a risk manager is able to target the most significant sources of risk, or the portfolio's so-called 'Hot Spots'. This task is complicated by the fact that VaR is a sub-additive measure: the portfolio VaR is typically less than the sum of the individual position VaRs. However, since VaR is a homogeneous function (i.e., $VaR(ax) = a \cdot VaR(x)$), it admits a marginal decomposition. Note that if we multiply Equation 4 by the position and sum over all holdings in the portfolio, we obtain

$$\sum_{i=1}^{N} x_i \frac{\partial VaR(x)}{\partial x_i} = VaR(x) \tag{5}$$

In Equation 5, each term in the summation is the product of position size and the rate of change of VaR with respect to that position. This essentially represents the rate of change of VaR with respect to a small *percentage* change in the size of the position. Let us define

$$C(x_i) = \frac{1}{VaR(x)} \times x_i \frac{\partial VaR(x)}{\partial x_i} \times 100\% \tag{6}$$

to be the percentage contribution to VaR of the i-th position. Equation 6 must be interpreted on a marginal basis; it indicates the relative contributions to the change in VaR that results if all positions are scaled by the same amount. Note that at the best hedge position, a position's marginal VaR, and therefore also its VaR contribution, is zero.

Similarly, multiplying both sides of Equation 3 by $m(x)^T$ shows that VaR is equal to the inner product of the VaR map and the VaR gradient with respect to the risk factor exposures. We can therefore define

$$C(m_k(x)) = \frac{1}{VaR(x)} \times m_k(x) \frac{\partial VaR(x)}{\partial m_k(x)} \times 100\% \tag{7}$$

to be the percentage contribution to VaR of the k-th risk factor. Again, Equation 7 must be interpreted on a marginal basis.

An Example FX Portfolio

Table 1 shows a portfolio of foreign exchange (FX) forward contracts as of July 1, 1997. Suppose that the exchange rates, in USD, are 0.73 (CAD), 0.58 (DEM), 0.17 (FRF) and 0.0090 (JPY). The total value of the portfolio is 122,000 USD and its one-day 99% VaR is 78,000 USD.

For this example, we elect to use the RiskMetrics risk factor data set for computing the parametric VaR. The portfolio's VaR map, along with the marginal VaR and VaR contribution of each risk factor, are shown in Table 2.

The risk factors in Table 2 are sorted in order of decreasing VaR contribution. Since the forwards are contracts to purchase foreign currency with USD, the VaR map consists of long positions in FX spots and foreign interest rates, and short positions in US interest rates. The magnitudes of these positions indicate that the portfolio has significant exposure to the DEM/USD exchange rate and 30- and 90-day interest rates in the US and Germany. Note that exposure to any risk factor can be eliminated by undoing the corresponding VaR map position (e.g., Canadian currency risk can be removed by selling 1.09 million USD worth of Canadian dollars).

DEM/USD exchange	6,327	102.67	82.99
FRF/USD exchange	1,355	95.62	16.55
JPY/USD exchange	90	97.51	1.12
US 90-day rate	-4,146	-0.10	-0.06
Germany 90-day rate	3,698	0.04	0.02
Germany 180-day rate	602	0.21	0.02
Canada 30-day rate	728	0.13	0.01
US 180-day rate	-628	-0.08	0.01
US 30-day rate	-5,631	0.01	0.00
Germany 30-day rate	3,517	-0.01	0.00
Canada 90-day rate	358	0.09	0.00
France 90-day rate	98	-0.21	0.00
Japan 30-day rate	33	0.01	0.00
Canada 180-day rate	20	-0.83	0.00
France 30-day rate	1,511	0.03	-0.01
CAD/USD exchange	1,090	-5.41	-0.75

Table 2: Risk factor data for the FX portfolio (ranked by VaR contribution)

The primary source of portfolio risk is currently the DEM exposure, as indicated by the fact that it contributes 83% of the VaR. Conversely, the CAD exposure is actually acting as a hedge for the portfolio, as evidenced by its negative VaR contribution (-0.75%). Furthermore, the VaR gradient indicates that, at the margin, the VaR can be reduced by 5.41 USD for every additional 10,000 USD of exposure to the Canadian dollar. This may be somewhat surprising given that the portfolio is currently long CAD; it is a useful illustration of the fact that portfolio risk depends not only on the individual risk factors themselves, but also on their correlation (in this case, the Canadian dollar is negatively correlated with the other currencies).

The previous analysis considers the effects of general market factors on the VaR. The analysis at the position level (Table 3) yields consistent results. Together, the two DEM contracts contribute approximately 83% of

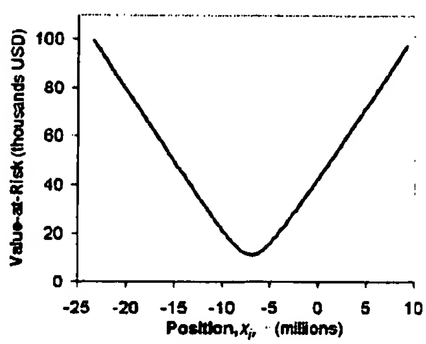

Figure 2: Trade risk profile for DEM/USD .57 60d

the current portfolio risk while the CAD contracts, with a total contribution of –0.70%, act as a hedge. Note that these values agree well with the VaR contributions of the DEM (83%) and CAD (–0.75%) risk factor exposures. The marginal VaRs indicate that increasing the positions in the DEM, FRF and JPY contracts results in greater portfolio risk while a similar increase in the CAD contracts reduces risk. This is also reflected by the best hedge positions; for the DEM, FRF and JPY contracts, the best hedges are smaller than the current positions (and in fact suggest shorting contracts), but they are larger than the current positions in the case of the CAD contracts.

The impact of holding the best hedge position in a given instrument can be measured in terms of the percentage reduction in VaR that can be achieved (i.e., the resulting decrease in VaR expressed as a percentage of the current VaR). At their best hedge positions, the DEM contracts each reduce the VaR by almost 88% while each CAD contract offers a much smaller reduction of only 0.2%. In many cases, however, it may simply not be feasible to hold an instrument at its best hedge position. For example, being short seven million units of DEM/USD .57 60d contracts may well run counter to the underlying objectives of the portfolio. Thus, it is often useful to consult the trade risk profile (e.g., Figure 2) to determine the VaR reduction that can be achieved within practical limitations.

Simulation-based VaR

The simulation-based approach to VaR calculation relies on a complete valuation of the portfolio under a set of scenarios, which may derive from historical data or a Monte Carlo simulation. Given a particular 'base case' scenario (e.g., representative of current market conditions), it is straightforward to calculate the gain or loss in portfolio value in each scenario. Let v_i^0 denote the unit value of instrument i in the base scenario and v_{ij}^t denote its unit value in scenario j at some future time t. We refer to v_{ij}^t as a **Mark-to-Future** value for instrument i. Since we assume exclusively a one-day time horizon for calculating VaR, we will hereafter dispense with the t superscript to improve readability. Let us define $\Delta v_{ij} = v_i^0 - v_{ij}$ to be the unit loss of instrument i in scenario j. If the current position in instrument i is x_i, then the loss (note that a gain is a negative loss) incurred by the portfolio in scenario j is

(8) $$L_j(x) = \sum_{i=1}^{N} x_i \, \Delta v_{ij}$$

Suppose that the likelihood, or weight, of scenario j is p_j. If we order the losses from largest to smallest (since losses can be negative when the portfolio gains in value, 'smallest' is taken here to mean 'most negative') and calculate the cumulative scenario probability, then the non-parametric $100(1-\alpha)$% VaR, or nVaR, equals the loss in that scenario for which the cumulative probability first meets or exceeds α. We refer to this scenario as the **threshold scenario**. To simplify the notation, we denote the threshold scenario simply as s^α, implicitly recognizing its dependence on x and α.

For example, consider a portfolio that is evaluated over a set of 100 scenarios. Table 4 shows the five largest losses, in decreasing order of magnitude, along with their

DEM/USD .57 60d	59.24	45.4	6.0	-7.0	87.7
DEM/USD .59 120d	58.97	37.7	5.0	-8.0	87.6
FRF/USD .16 40d	16.19	16.5	8.0	-38.3	79.2
JPY/USD .0091 11d	0.88	1.1	10.0	-209.2	13.2
CAD/USD .73 100d	-3.80	-0.2	0.5	1.4	0.2
CAD/USD .74 30d	-3.85	-0.5	1.0	1.9	0.2

Table 3: Instrument data for the FX portfolio (ranked by VaR contribution)

respective scenarios, probabilities and cumulative probabilities. For this particular portfolio and scenario set, the 95% and 98% nVaRs are 8,800 and 9,500, respectively. Note that in the latter case, $p_{27} + p_{82} = 0.04 > 0.02$ and so one might argue that some value between 10,000 and 9,500 provides a better estimate for the 98% nVaR. A discussion of the merits of such interpolation schemes is beyond the scope of this paper; we simply note that the above approach may yield a smaller nVaR, relative to interpolated values, in some cases.

27	10,000	0.010	0.010
82	9,500	0.030	0.040
50	8,800	0.010	0.050
11	8,600	0.020	0.070
63	8,100	0.005	0.075

Table 4: Simulation-based VaR example

While valuing the portfolio under large numbers of scenarios can be a computationally intensive task, nVaR analysis has the desirable property of requiring only a single simulation. Once the instruments' Mark-to-Future values have been obtained, Equation 8 can be used to calculate losses for individual holdings (and hence the portfolio) under subsequent changes in the positions.

Trade Risk Profile and Best Hedge Position

Recall that the trade risk profile plots the level of risk (VaR or nVaR) against the position taken in a particular instrument. In the parametric case, the resulting curve is smooth and has a unique minimum at the best hedge position. The nVaR's dependency on a finite number of scenarios implies that the non-parametric trade risk profile (nTRP) is piecewise linear (Figure 3). As shown in the Appendix, the nTRP consists of multiple segments, each corresponding to a threshold scenario that is in effect for a given range of positions. Unlike the parametric case, the nTRP may have multiple local minima, and therefore, finding the best hedge position (i.e., the global minimum), x_i^{nbh}, may require considerable computational effort. However, this effort lies only in tracing out the nTRP, rather than re-pricing the instruments, and so computational time is typically far less than that required for a full simulation.

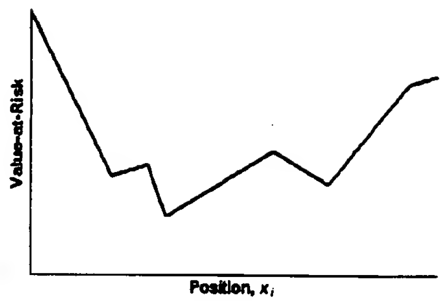

Figure 3: Simulation-based trade risk profile

Marginal nVaR

One might anticipate that calculating the marginal nVaR involves making a small positional change, re-simulating the portfolio and recalculating nVaR. Fortunately, this is not required. From the definition of nVaR as the loss in the threshold scenario

$$nVaR(x) = \sum_{i=1}^{N} x_i \, \Delta v_{is}, \tag{9}$$

it follows that nVaR is linear in x. Let us assume for the moment that the threshold scenario remains unchanged for small variations in the positions. In this case, the derivative of nVaR with respect to the i-th position is

$$\frac{\partial \, nVaR(x)}{\partial x_i} = \Delta v_{is}. \tag{10}$$

Thus, the i-th component of the nVaR gradient is simply the difference between the instrument's values in the base and threshold scenarios.

Let us now examine the marginal nVaR in light of the piecewise linearity of the nTRP. In doing so, we will make reference to Figure 4, which illustrates two adjacent segments of a nTRP for some instrument i. In this case, positions (x_i) of 100, 200 and 300 result in portfolio nVaR values of 25,000, 10,000 and 15,000, respectively. The slope of the first segment is −150 while that of the second segment is 50. Recall that positions in all instruments other than i are held fixed.

Since the nTRP is piecewise linear, all positions in instrument i that lie on the same segment have an identical gradient, whose i-th component is simply the slope of that segment; that is, for all $100 < x_i < 200$, increasing (decreasing) the position in instrument i by a sufficiently small amount δ decreases (increases) the portfolio nVaR by 150δ. In this case, a 'sufficiently small' change in position can be calculated precisely as a decrease of up to $x_i - 100$ or an increase of up to $200 - x_i$.

Now consider a position of 200 in instrument i, which corresponds to a change in the threshold scenario. As is evident in Figure 4, the gradient is not well-defined at this point; its i-th component changes abruptly from -150 to 50. To deal with this lack of continuity in the gradient at such points, it is necessary to consider two one-sided sub-gradients. It should be apparent, however, that knowledge of the slopes and endpoints of the segments comprising the nTRP allows marginal nVaR information, and the range for which it is valid, to be reported for any position.

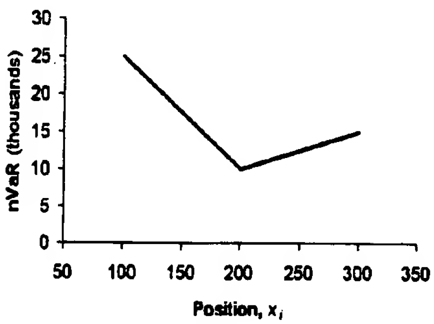

Figure 4: Two segments of a nTRP

nVaR Contribution

Equation 9 immediately provides us with a risk decomposition—nVaR is the sum of the position losses in the threshold scenario. Thus, the percentage contribution to nVaR of the i-th position is simply

$$nC(x_i) = \frac{1}{nVaR(x)} \times x_i \Delta v_{is^*} \times 100\% \qquad (11)$$

Note that Equation 11 is identical to Equation 6 in that the risk contribution is based on the product of the position and the marginal nVaR. Thus, as in the parametric case, the above decomposition must be interpreted on a marginal basis. If we scale all positions by some factor $(1 + \varepsilon)$, where ε is a small constant, then nVaR increases by an amount $\varepsilon \times nVaR(x)$ and Equation 11 indicates the relative contribution of the i-th instrument to this increase.

Implementation Considerations

The simulation-based approach to VaR, as described in this paper, depends entirely on the scenarios used in the simulation; changing the scenarios is likely to yield different values for nVaR as well as for the related marginal risk measures. Thus, it is important to recognize the possible effects of sampling error on the reported values. In particular, while increasing the number of scenarios generally improves the reliability of nVaR as an estimate of the true Value-at-Risk, the marginal nVaR (and similarly, the nVaR contribution) may still exhibit considerable variability as more scenarios are sampled. To improve the accuracy of these values, we propose using a smooth approximation to the nTRP.

The Problem: Sensitivity to the Threshold Scenario

Recall from Equation 10 that the marginal nVaR is determined exclusively by an instrument's values in the base and threshold scenarios. Obtaining consistent marginal nVaR estimates, then, requires that scenarios resulting in similar losses also have similar Mark-to-Future values for the instruments. However, this may not hold in practice, as illustrated by the following example. Consider a portfolio consisting of only two positions that is simulated over two scenarios (Table 5). The portfolio has an identical value (and therefore an identical loss) in both scenarios, yet the instruments' Mark-to-Future values (and the marginal nVaRs) are quite different. Therefore, the marginal nVaR is extremely sensitive to the threshold scenario.

[redacted]

1	100	10	5
2	50	10	20
Portfolio		1,500	1,500

Table 5: Sample two-instrument portfolio

Since an instrument's marginal nVaR equals the slope of the nTRP for that instrument at the current position, it follows that adjacent segments of the nTRP can have markedly different slopes (recall from Figure 4 that adjacent segments of the nTRP meet at points where the respective threshold scenarios incur the same loss). Increasing the number of scenarios tends to shorten the average length of segments comprising the nTRP.

However, adjacent segments do not necessarily become better 'aligned' in the sense of having similar slopes. Thus, the marginal nVaR may not exhibit the convergence that might be expected as more scenarios are sampled.

The Solution: Smoothing the nTRP

Fitting a smooth curve to the nTRP, and then taking the derivative of this curve, tends to provide a more robust estimate of the true marginal VaR. Essentially, this approach removes the 'noise' that is present in the nTRP.

One might consider using splines or fitting a polynomial function, $P_i(x_i)$, to the nTRP of instrument i (specifically, to the endpoints of the segments) in the least squares sense. Clearly, the degree of $P_i(x_i)$ should be chosen so that the curve provides a reasonable approximation to the nTRP, without over-fitting the points. A visual comparison of the two curves will generally establish the suitability of $P_i(x_i)$ for the range of positions being considered. From this approximation, one can then obtain the following estimates:

· marginal nVaR at position x_i

$$\frac{\partial\, nVaR(x)}{\partial x_i} = P_i'(x_i)$$

· best hedge position

$$x_i^{nbh} = \left\{ x_i^* \mid P_i(x_i^*) \le P_i(x_i) \text{ for all } x_i \text{ in the range of interest} \right\}$$

· nVaR contribution at position x

(12) $$nC(x_i) = \frac{x_i\, P_i'(x_i)}{\sum_{l=1}^{N} x_l\, P_l'(x_l)} \times 100\%$$

The FX Portfolio Revisited

To compare the parametric and simulation-based VaR analyses, the FX portfolio was simulated over a set of 1,000 Monte Carlo scenarios. The one-day 99% nVaR is 77,000 USD, which differs from the parametric value by less than 2%. The resulting loss histogram is approximated well by a normal distribution (Figure 5) and the nVaR analysis (Table 6) is consistent with its parametric counterpart in Table 3. The final column of Table 6 gives the range for which the marginal nVaR remains unchanged (i.e., the 'length' of the current segment of the nTRP for each instrument). Note the close agreement between the parametric and simulation-based trade risk profiles for DEM/USD .57 60d (Figure 6).

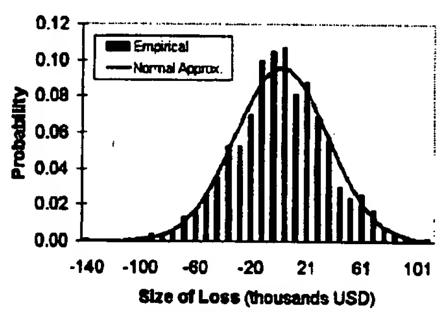

Figure 5: Distribution of losses for the FX portfolio with best normal approximation (1,000 scenarios)

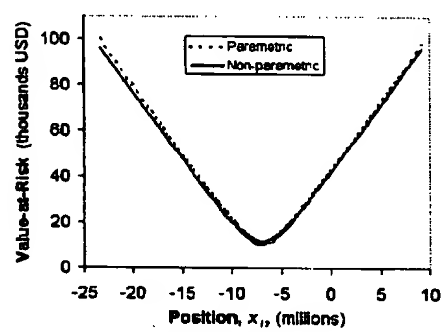

Figure 6: TRP and nTRP for DEM/USD .57 60d

Example: the NIKKEI Portfolio

Table 7 shows a portfolio that implements a butterfly spread on the NIKKEI index, as of July 1, 1997. In addition to common shares of Komatsu (current price 840,000 JPY) and Mitsubishi (current price 860,000 JPY), the portfolio includes several European call and put options on these equities. The total value of the portfolio is 12,493 million JPY and its parametric one-day 99% VaR is 115 million JPY.

This portfolio, which may be representative of the positions held by a trading desk, makes extensive use of options to achieve the desired payoff profile. A histogram showing the distribution of losses over a set of 1,000 Monte Carlo scenarios (Figure 7) indicates that the normal distribution fits the data poorly, and that the parametric VaR is likely to over-estimate the true Value-at-Risk. Indeed, simulating the portfolio over these 1,000 scenarios results in a one-day 99% nVaR of 2.9 million JPY, reflecting the fact that the portfolio is well-hedged. Because the parametric VaR measures the risk poorly in this case, we perform only a simulation-based analysis.

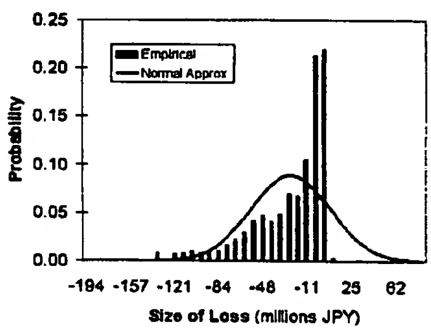

Figure 7: Distribution of losses for the NIKKEI portfolio with best normal approximation (1,000 scenarios)

We compute the contributions, marginal nVaRs and best hedge positions using a pure non-parametric approach (i.e., corresponding to the piecewise-linear nTRP), as well as a third-order polynomial approximation to the nTRP. This analysis is summarized in Table 8. The results of the polynomial approximation appear in parentheses.

The magnitudes of the nVaR contributions are quite large, ranging from –2387% (Mitsubishi Psep30 800) to 2151% (Komatsu Cjun2 670). This is due to the fact that the portfolio is highly-leveraged and well-hedged, so that the risks incurred by individual positions tend to offset each other to a large extent. In particular, considering the relative sizes of the positions in the portfolio, note that the Komatsu Cjun2 670 position stands to gain considerably if the market appreciates while the Mitsubishi Psep30 800 position acts in the opposite manner. More generally, as indicated by their negative contributions, the two short calls and the two long puts act as a hedge for the portfolio, protecting against drops in the NIKKEI index.

The contributions (and the marginal nVaRs) calculated using the polynomial approximation to the nTRP are roughly half the size of those obtained by the pure non-parametric approach. As will be discussed shortly, this is due to the smoothing effects of the approximation. Note, however, that the relative sizes of the contributions among all instruments are the same in both cases (i.e., they yield an identical ranking of the instruments in terms of the nVaR contribution).

Based on the marginal nVaR, the most attractive possibilities for lowering the overall portfolio risk include any one of the following trades: reducing the current holdings in Komatsu Cjun2 670 or Komatsu Cjun2 760, selling one of the common stocks, or shorting additional calls on Mitsubishi (i.e., Mitsubishi Cjul29 800). Purchasing additional units of Mitsubishi Psep30 800 is also a promising option.

If it is feasible to hold an instrument at its best hedge position, then Komatsu Cjun2 760 offers the greatest potential for reducing risk (i.e., a reduction of 42.2%).

DEM/USD .57 60d	59.73	46.6	6.0	–7.4	87.0	[2.7, 9.2]
DEM/USD .59 120d	59.62	38.7	5.0	–8.2	86.9	[1.7, 8.2]
FRF/USD .16 40d	15.03	15.6	8.0	–38.0	78.6	[3.4, 20.8]
JPY/USD .0091 11d	1.01	1.3	10.0	–258.6	18.8	[–3.1, 111.1]
CAD/USD .73 100d	-11.86	–0.8	0.5	1.0	0.8	[–2.5, 1.0]
CAD/USD .74 30d	-11.01	–1.4	1.0	1.6	0.8	[–2.3, 1.6]

Table 6: nVaR analysis for the FX portfolio (1,000 scenarios)

Komatsu	Equity	n/a	n/a	2.5	2,100,000
Mitsubishi	Equity	n/a	n/a	2.0	1,720,000
Komatsu Cjul29 900	Call	7	900	-28.0	-11,593
Mitsubishi Cjul29 800	Call	7	800	-16.0	-967,280
Mitsubishi Csep30 836	Call	70	836	8.0	382,070
Mitsubishi EC 6mo 860	Call	184	860	11.5	563,340
Komatsu Cjun2 760	Call	316	760	7.5	1,020,110
Komatsu Cjun2 670	Call	316	670	22.5	5,150,461
Komatsu Paug31 760	Put	40	760	-10.0	-68,919
Komatsu Paug31 830	Put	40	830	10.0	187,167
Mitsubishi Psep30 800	Put	70	800	40.0	2,418,012

Table 7: NIKKEI portfolio

Note the close agreement between the best hedge positions as determined by the nTRP and by the polynomial approximation.

Smoothing and the Polynomial Approximation

Figure 8 shows the nTRP and its polynomial approximation for one of the Mitsubishi call options. At the current position (11,500), the nTRP is more steeply sloped than the polynomial, which results in a larger marginal nVaR (i.e., 575 versus 329). Note that smoothing counteracts discrepancies caused by the piecewise linearity of the nTRP. This is particularly evident at a position of 7,000; here, the nTRP slopes upwards, implying a positive marginal nVaR, while the negative marginal nVaR derived from the approximation is more consistent with the general shape of the trade risk profile.

Komatsu Cjun2 670	2161	2727	22.5	20.1	41.3
	(1094)	(1028)		(20.0)	(43.0)
Komatsu Cjun2 760	678	2576	7.5	5.0	42.2
	(344)	(970)		(4.8)	(44.2)
Mitsubishi Csep30 836	477	1699	8.0	3.9	34.8
	(247)	(653)		(4.6)	(35.3)
Mitsubishi EC 6mo 860	232	575	11.5	9.1	22.9
	(179)	(329)		(7.8)	(23.0)
Komatsu	202	2300	2.5	-0.4	35.1
	(101)	(857)		(-0.1)	(35.0)
Mitsubishi	149	2119	2.0	-1.2	35.1
	(75)	(790)		(-0.8)	(35.0)
Komatsu Paug31 760	53	-152	-10.0	33.2	35.2
	(28)	(-60)		(28.7)	(35.7)
Komatsu Cjul29 900	-51	52	-28.0	150.5	34.7
	(-26)	(20)		(-121.0)	(35.2)
Komatsu Paug31 830	-237	-675	10.0	19.8	34.6
	(-119)	(-252)		(19.0)	(35.4)
Mitsubishi Cjul29 800	-1166	2078	-16.0	-19.3	34.8
	(-591)	(781)		(-18.8)	(35.0)
Mitsubishi Psep30 800	-2387	-1702	40.0	44.2	34.7
	(-1232)	(-651)		(43.5)	(35.4)

Table 8: Analysis of the NIKKEI portfolio based on 1,000 scenarios (results of polynominal approximation in parenthesis)

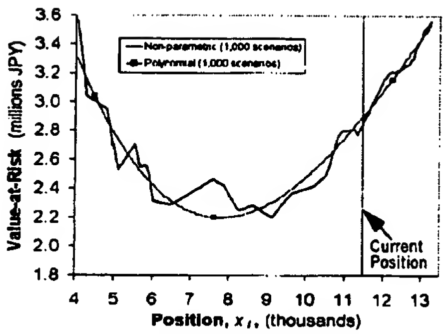

Figure 8: nTRP and polynomial approximation for Mitsubishi EC 6mo 860

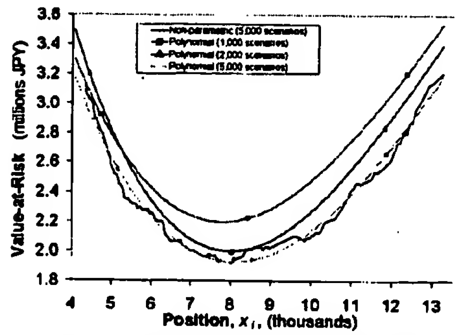

Figure 9: nTRP and polynomial approximations for Mitsubishi EC 6mo 860

To illustrate the effects of increasing the number of scenarios on the nTRP, the NIKKEI portfolio was also simulated over 2,000 and 5,000 scenarios. These scenario sets are obtained by first adding 1,000, and then a further 3,000 scenarios to the initial scenario set. Figure 9 plots the 5,000-scenario nTRP along with each of the polynomial approximations. The curves tend to shift downwards, suggesting a smaller nVaR, as the number of scenarios is increased. However, we note that they remain within the 95% confidence interval for nVaR, calculated for the 1,000-scenario simulation, at the current (11,500) and best hedge positions (9,100). While increasing the number of scenarios creates more segments in the nTRP (compare Figures 8 and 9), sampling error remains a concern even at the 5,000-scenario level (i.e., one can find segments on the nTRP whose slopes are inconsistent with those of the polynomial approximation). In contrast, the approximations tend to provide more consistent gradient information.

It should be noted, however, that even when using the polynomial approximation, we observe inconsistencies in the nVaR contributions (i.e., a position that contributes positively to VaR in one analysis is found to have a negative contribution in another). Since the portfolio nVaR is extremely small relative to the individual position nVaRs in this case, slight errors in the polynomial approximations may combine to change the sign of the denominator in Equation 12. This should only be viewed as a concern for well-hedged, highly-leveraged portfolios, such as the one considered in this example, rather than for portfolios containing options in general.

Conclusions

This paper has examined tools for VaR-based risk management. Tools for decomposing VaR, assessing its marginal impacts and constructing best hedges, allow managers to understand the sources of risk better and to manipulate the portfolio to effect the desired changes in risk. The analytical techniques that derive from the parametric, or delta-normal, VaR form the basis of a risk manager's toolkit when dealing with portfolios of linear instruments. The simulation-based tools developed in this paper extend these capabilities to portfolios that contain non-linearities or are subject to non-normal market distributions. An attractive feature of these methods is their need for only a single simulation to obtain the Mark-to-Future values of the instruments. Furthermore, it is straightforward to incorporate new instruments into the analysis by simulating them independently of the portfolio itself. Thus, while our analyses considered trading only those instruments currently held in the portfolio, it extends naturally to encompass so-called incremental VaR. Specifically, it is only necessary to simulate the additional instruments to be considered. Once their Mark-to-Future values have been obtained, the instruments can be easily incorporated in any marginal nVaR analysis by assigning them a current position of zero.

We note that sampling errors can occasionally yield inconsistent results under the simulation-based approach to VaR. Hence, we propose fitting a smooth curve to the (piecewise-linear) trade risk profile to obtain more robust estimates of the VaR contribution, marginal VaR and best hedge positions. The techniques are demonstrated on a portfolio of European options that is poorly-suited for parametric analysis. The results show that a smooth approximation to the nTRP improves the reliability of

marginal VaR estimates, although caution is required when interpreting VaR contributions for well-hedged, highly-leveraged portfolios.

The accuracy of the simulation-based analysis, in light of potential sampling error, is a subject worthy of further investigation. The polynomial approximation to the nTRP is a fairly simple one; the use of splines or other functions may be a preferable approach. In this paper, we have only considered smoothing the trade risk profiles. A more ambitious strategy might seek to fit a probability distribution to the losses themselves and then estimate VaR based on this distribution. This would in fact eliminate the piecewise linearity of the nTRP that is characteristic of the current approach. Furthermore, we anticipate that alternative risk measures, such as expected shortfall, or regret, will exhibit more robust behaviour than nVaR, which relies exclusively on the threshold scenario. ◎

Helmut Mausser joined Algorithmics in November 1997 as the recipient of an NSERC Industrial Research Fellowship. He is a mathematician in the Quantitative Research group. Helmut has a Ph.D. in Operations Research from the University of Colorado.

Dan Rosen joined Algorithmics in 1995. He is currently the director of Research and is responsible for Algorithmics' financial and mathematical research and joint projects with academic institutions. Dan has a Ph.D. in Applied Sciences from the Department of Chemical Engineering and Applied Chemistry at the University of Toronto.

References

Artzner, P., F. Delbaen, J.M. Eber and D. Heath, 1998,
"Coherent measures of risk," Working Paper, Institut de Recherche Mathématique Avancée, Université Louis Pasteur et C.N.R.S.

Dowd, K., 1998,
"VaR by increments," Enterprise Wide Risk Management Special Report, *Risk*: 31-32.

Garman, M., 1996,
"Improving on VaR," *Risk* 9(5): 61-63.

Garman, M., 1997,
"Taking VaR to pieces," *Risk* 10(10): 70-71.

Litterman, R., 1996a,
"Hot spots and hedges," Risk Management Series, Goldman Sachs.

Litterman, R., 1996b,
"Hot spots and hedges," *Journal of Portfolio Management* (Special Issue): 52-75.

Litterman, R., 1997a,
"Hot spots and hedges (I)," *Risk* 10(3): 42-45.

Litterman, R., 1997b,
"Hot spots and hedges (II)," *Risk* 10(5): 38-42.

Markowitz, H., 1952,
"Portfolio selection," *Journal of Finance* 7(1): 77-91.

Mausser, H. and D. Rosen, 1998,
"Marginal VaR analysis: parametric and non-parametric methods," Research Paper, Algorithmics Inc.

RiskMetrics™ Technical Document, 4th Edition, J.P. Morgan, December 1996.

Sharpe, W., 1964,
"Capital asset prices: a theory of market equilibrium under conditions of risk," *Journal of Finance* 19(3): 425-442.

Appendix

Calculating Parametric VaR

Consider a portfolio with N holdings that is exposed to W market risk factors. Each instrument in the portfolio is decomposed into a set of risk factor positions so that the change in the instrument's value, Δv_i, can be expressed linearly in terms of the risk factor returns:

$$\text{(A1)} \qquad \Delta v_i = \sum_{k=1}^{W} m_k^i r_k$$

Recall that the vector m^i is the VaR map of instrument i. We can express the change in the value of the portfolio as the sum of the changes in the values of its holdings:

$$\Delta V(x) = \sum_{i=1}^{N} x_i \sum_{k=1}^{W} m_k^i r_k$$

From the definition of the portfolio VaR map (Equation 2), we can write Equation A1 more compactly as

$$\Delta V(x) = m(x)^T r$$

Note that $\Delta V(x)$ is normally distributed with mean zero and variance $m(x)^T Q^* m(x)$, so that the $100(1-\alpha)$% VaR is

$$VaR(x) = Z_\alpha \sqrt{m(x)^T Q^* m(x)}$$

Defining $Q = Z_\alpha^2 Q^*$ (e.g., $Z_{.05} = 1.645$) yields Equation 1.

Parametric Trade Risk Profile

To construct the trade risk profile for instrument i, fix the positions in all instruments other than i to their current values and consolidate them into a base portfolio position x_l. Denote the single-unit volatilities of instrument i and the base portfolio by σ_i and σ_l, respectively, and their correlation by ρ_{il}. The volatility of the portfolio is

$$\sigma(x) = \sqrt{(x_i\sigma_i)^2 + (x_l\sigma_l)^2 + 2\rho_{il}(x_i\sigma_i)(x_l\sigma_l)}$$

Since

$$VaR(x) = Z_\alpha \sigma(x)$$

it follows that the trade risk profile is a curve of the form

$$\text{(A2)} \qquad f(x_i) = \sqrt{a x_i^2 + b x_i + c}$$

where $a = (Z_\alpha \sigma_i)^2$, $b = 2Z_\alpha^2 \rho_{il}\sigma_i\sigma_l x_l$ and $c = (Z_\alpha \sigma_l x_l)^2$. Differentiating Equation A2 with respect to x_i yields

$$\frac{d f(x_i)}{dx_i} = \frac{2a x_i + b}{2 f(x_i)}$$

Since $f(x_i)$ is strictly positive, the unique minimum occurs at the best hedge position

$$x_i^* = -\frac{b}{2a} = -\frac{\rho_{il}\sigma_l x_l}{\sigma_i}$$

It is straightforward to show that $f(x_i)$ is symmetric around this point.

Simulation-based Trade Risk Profile

To construct the nTRP for instrument i, let us fix the positions in all instruments other than i to their current values. The loss incurred by the portfolio in scenario j (see Equation 8) can be written

$$L_j(x_i) = \Delta V_{ij} + x_i \Delta v_{ij}$$

where ΔV_{ij} includes the losses due to all instruments other than i. If we plot portfolio losses against the position in instrument i, then each scenario j gives rise to a line $L_j(x_i)$ with slope Δv_{ij} and x-intercept $-\Delta V_{ij} / \Delta v_{ij}$. The piecewise linearity of the nTRP follows from the fact that it is composed of segments from these lines. Specifically, the nTRP consists of the segments defined by the threshold scenario at each position. This is illustrated in Figure A1, which shows a nTRP in which the threshold scenario is always the one with the third-largest loss. Note that each nTRP segment lies on the third line from the top.

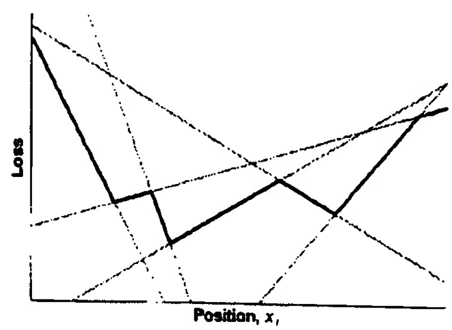

Figure A1: Example of a nTRP (in bold)

In general, the threshold scenario can change whenever the line $L_j(x_i)$ intersects that of another scenario. An algorithm that constructs the trade risk profile and finds the best hedge position is described in Mausser and Rosen (1998).